

## Article

# The SUMO Ship Detector Algorithm for Satellite Radar Images

Harm Greidanus \*, Marlene Alvarez, Carlos Santamaria, François-Xavier Thoorens, Naouma Kourti and Pietro Argentieri

European Commission, Joint Research Centre (JRC), Directorate for Space, Security & Migration, Via E. Fermi 2749, I-21027 Ispra (VA), Italy; marlene.alvarez@ec.europa.eu (M.A.); carlos.santamaria@ec.europa.eu (C.S.); fx.thoorens@ark.io (F.-X.T.); naouma.kourti@ec.europa.eu (N.K.); pietro.argentieri@ext.ec.europa.eu (P.A.)

\* Correspondence: harm.greidanus@ec.europa.eu; Tel.: +39-0332-789-739

Academic Editors: Xiaofeng Li and Prasad S. Thenkabail

Received: 25 November 2016; Accepted: 13 February 2017; Published: 7 March 2017

**Abstract:** Search for Unidentified Maritime Objects (SUMO) is an algorithm for ship detection in satellite Synthetic Aperture Radar (SAR) images. It has been developed over the course of more than 15 years, using a large amount of SAR images from almost all available SAR satellites operating in L-, C- and X-band. As validated by benchmark tests, it performs very well on a wide range of SAR image modes (from Spotlight to ScanSAR) and resolutions (from 1–100 m) and for all types and sizes of ships, within the physical limits imposed by the radar imaging. This paper describes, in detail, the algorithmic approach in all of the steps of the ship detection: land masking, clutter estimation, detection thresholding, target clustering, ship attribute estimation and false alarm suppression. SUMO is a pixel-based CFAR (Constant False Alarm Rate) detector for multi-look radar images. It assumes a K distribution for the sea clutter, corrected however for deviations of the actual sea clutter from this distribution, implementing a fast and robust method for the clutter background estimation. The clustering of detected pixels into targets (ships) uses several thresholds to deal with the typically irregular distribution of the radar backscatter over a ship. In a multi-polarization image, the different channels are fused. Azimuth ambiguities, a common source of false alarms in ship detection, are removed. A reliability indicator is computed for each target. In post-processing, using the results of a series of images, additional false alarms from recurrent (fixed) targets including range ambiguities are also removed. SUMO can run in semi-automatic mode, where an operator can verify each detected target. It can also run in fully automatic mode, where batches of over 10,000 images have successfully been processed in less than two hours. The number of satellite SAR systems keeps increasing, as does their application to maritime surveillance. The open data policy of the EU's Copernicus program, which includes the Sentinel-1 satellite, has hugely increased the availability of SAR images. This paper aims to cater to the consequently expected wider demand for knowledge about SAR ship detectors.

**Keywords:** ship detection; maritime surveillance; satellite images; radar images; synthetic aperture radar

## 1. Introduction

Many parties, from both public and private sectors, have the need to be aware of the presence of ships and shipping traffic at sea. A number of systems can provide data on this, and a useful distinction can be made between cooperative/reporting systems, where ships provide data themselves, and non-cooperative/observation systems, where sensors are used to obtain data not relying on the ships' cooperation. Among the latter systems, imaging radar deployed on orbiting satellites is

a valuable type of sensor, which can collect data worldwide. A radar, in order to make its image, illuminates its scene of observation with radar waves, and ships reflect those waves back to the radar, producing localized bright spots in the image [1,2]. The surrounding sea surface also scatters a fraction of the incoming radar waves back, and physical features on the sea, such as waves, fronts, reefs or rain, can produce increased and localized radar backscatter [3,4]. Rain, in fact, can result in increased or in decreased backscatter, depending on the radar wavelength [5,6], but its main impact on ship detection is from localized increases. In addition, the coherent nature of radar causes speckle noise [1]. All of this leads to a noisy/cluttered background in the maritime radar image. Ship detection is the task of finding in the image the reflections of ships in the background of the sea clutter [7].

Various types of radars are operated on satellites, the main ones being scatterometer, altimeter, real aperture imaging radar and Synthetic Aperture Radar (SAR). Almost all imaging radars on satellites are of the SAR type, because only that technique leads to a high resolution at the long ranges that characterize space-based observation [1]. Recent years have seen an accelerating increase in the number of imaging radar satellites in orbit. After the gap that followed the short period of operation of Seasat in 1978, satellite SAR imagery has been continuously available since the launch of ERS-1 in 1991. At the time of writing, about 10 SAR satellites (some part of a constellation) are active. While none of them have been uniquely designed or optimized for ship detection (the driving applications are land monitoring and sea ice monitoring), all are suitable for that task, and some have imaging modes that are very suitable for, or even (within the design constraints) dedicated to, maritime surveillance.

Many satellite SARs allow a choice of imaging modes, ranging between wide area coverage at low resolution to small area coverage at high resolution. In rough numbers, at the wide area end of the possibilities, the coverage can reach 500 km at a 70-m resolution, and at the other end, resolutions of 50 cm can be attained inside a 4 km coverage (the more attractive combination of the wide area at high resolution is not possible due to the fundamental limitations of radar imaging, as well as to limitations on the data load). The radar image is constructed as a collection of parallel lines that look away sideways from the radar's path of motion in what is called the range direction. Typical incidence angles (between the surface normal and the line of sight) along the range direction can be between  $20^\circ$  (steeply looking down) to  $60^\circ$  (looking far away to the side), but a particular image will only cover a sub-range of that. To form the 2D image, the lines are stacked along the radar's flight path, which is called the azimuth direction. Considering the sizes of ships, which range from 460 m for the largest ship down to practically human size, it is clear that the smallest ships will not be seen in wide-swath images, while small-to-medium ships in such images may show up as a small cluster of bright pixels. On the other hand, big ships can show detailed structure in very high-resolution images. The nature of the SAR imaging, however, leads to a ship signature in the radar image that does not correspond to the visual appearance of the ship. Instead, the ship's radar signature is dependent on details in the ship's structure down to scales of the radar wavelength and on the radar reflectivity of its materials. The signature is, in addition, quasi-random due to the sensitivity on the viewing geometry and the fading effect (interference of the coherent radar waves). Furthermore, motions of the ship cause a blurring in the SAR image. This behavior severely limits the amount of information that can be obtained from the SAR images of ships. Except for the highest resolution SAR images taken under good conditions, in most cases, the physical properties of a ship that can be extracted from its signature in the SAR image, besides geographic location, are limited to length, width and heading, and they often come with significant errors attached.

A single satellite SAR image contains of the order of  $10,000 \times 10,000$  pixels and is too big to fit on a normal computer screen. Although it is possible to scan the image visually for bright spots that represent ships, that would take much time. Apart from finding the ships, also the mentioned attributes (size and heading) need to be extracted. This means that a functionality (algorithm/software) is needed to automatically detect the ships and extract the ship attributes from a satellite radar image. Although conceptually it is not too difficult for automatic algorithms to find clusters of bright pixels down to a certain level, in reality, the inhomogeneity of the sea surface in the radar image makes it

challenging. Practice has shown that many of such detected clusters are not ships, but other objects or phenomena on the sea surface, either natural, such as small islands, reefs, rain cells, wind fronts or steep/breaking waves, or man-made, such as piers or port-related constructions. Automatic algorithms have difficulty with reliably discriminating between real ships and false alarms on such detections. Therefore, common practice has been that the automatic detection algorithm results are presented visually to a trained human operator, who is able to quickly discard many of the false alarms; although in many cases, even a human expert cannot be sure about the nature of a detection. Such a way of operating the ship detection algorithm is semi-automatic. Search for Unidentified Maritime Objects (SUMO) is designed with the choice to operate fully or semi-automatically.

Many algorithmic approaches to automatic ship detection in radar images have been explored in the literature. Useful reviews are given in [8,9]. For use on multi-look images, which contain only the amplitude (or power) of the radar backscatter, but not its phase, the most widely-used approach is the adaptive threshold or Constant False Alarm Rate (CFAR) detector, which looks for pixel values that are high compared to the local background (e.g., [10,11]). SUMO is of that type. Often, the local background of a pixel under test is estimated in a rectangular frame around it, excluding the center to prevent contamination of the background by the potentially present target. Sometimes not a single pixel, but the mean of a small box is compared to the local background [12]. (SUMO follows another approach as explained below.) Alternative approaches to the detection include wavelet analysis [13], use of a neural network [14,15], segmentation [16], non-parametric models [17] or saliency [18]. For use on complex SAR images, which SUMO does not exploit, yet other techniques are possible, such as split-look or sub-apertures [19], coherent-to-incoherent ratio [20] and polarimetric methods [21–25].

The purpose of this paper is to describe the algorithmic functionalities of the SUMO ship detector, which, according to the authors, has been representing the state-of-the-art for many years. The algorithms have been developed during a multi-year learning process based on the use of large amounts of real satellite SAR data. SUMO has been used in many scientific publications, in particular on fisheries' control [26–31], maritime surveillance [32–36] and benchmarking [37–40], but its detailed working has never been fully described.

## 2. Overview Description

### 2.1. Purpose of SUMO

SUMO is the name given to the ship detector. It is a set of algorithms, implemented in software, to find ships in satellite radar images in a semi-automatic or a fully automatic way. It was conceived of to enable the exploitation of satellite radar images for maritime surveillance with minimal human operator effort. Initial developments were aimed at fisheries' control, leading to the concept of VDS (Vessel Detection System) that employs satellite images to augment the VMS (Vessel Management System) self-reporting regime for fisheries [41]. Application areas were subsequently expanded to maritime safety and security. The SUMO software implements the set of algorithms that deal with the various sub-tasks of the ship detection process. Early versions of SUMO have been implemented in IDL (Interactive Data Language) and recent versions in Java; test versions have been implemented in MATLAB. However, the aim of this paper is to discuss the algorithm rather than any specific implementation.

### 2.2. General Approach

SUMO works on amplitude images, i.e., no complex (phase) information is used. The detection approach is fundamentally based on finding locally bright pixels on the sea. The sea clutter and speckle noise cause pixel values in the image, regardless of the presence of ships, to have a quasi-random distribution with variations around a local mean. The latter is mainly determined by local wind condition and incidence angle. A ship can only be discerned from the background if its radar reflection leads to pixel values that stand above the mean background plus noise. At the same time, some

background pixel values may accidentally attain similarly high values. The use of a threshold to select bright pixels to indicate the presence of ships unavoidably leads to the inclusion of false alarms from accidentally bright background pixels. This is a fundamental source of error for ship detection in radar images.

The high level functioning of SUMO consists of the following steps:

1. Ingestion of the satellite radar image and its metadata.
2. Selection of processing parameter values.
3. Application of a land mask; all subsequent processing applies only to the sea area.
4. Computation over the image of the local sea clutter level, i.e., the local background pixel statistics.
5. Derivation of a local detection threshold. All pixels brighter than the local threshold are “detected”. The threshold is computed, based on the local statistics and an assumed Probability Density Function (PDF) for the clutter, as a pixel value above which a clutter pixel has a fixed probability of occurring. This is the CFAR approach, and SUMO is a CFAR detector.
6. In the case of multi-polarization images, where the same scene is imaged in several polarization channels, Steps 4 and 5 are computed for each polarization channel separately, and all detected pixels are taken in union across the channels; i.e., a multi-channel pixel is detected if it is above the detection threshold in at least one channel.
7. All nearby detected pixels are clustered together into one detected object (from here on, “detection” will signify a detected object, whereas “detected pixel” will continue to be used to denote individual pixels above the CFAR threshold).
8. Extraction of the attributes of the detections: geographic location, length, width, heading, peak pixel value, integrated value, significance, and more. Length, width and heading are based on the notion that the target is an elongated cluster, and significance is a measure of how far the object sticks out from the clutter.
9. Discrimination of the detections between real ship and false alarm on the basis of the attributes and the assignment of a reliability value to the detection. The reliability value is calculated based on significance, length, width and the likelihood that the detection is an azimuth ambiguity (see below).
10. Optionally, for semi-automatic operation, inspection by a human operator of the detections in their surroundings, individually deciding on keeping/discarding.
11. Export of the results.

An early version of the algorithm [42] used a template matching approach, but that was abandoned for pixel-based thresholding, as the template-based approach was found to be not generically applicable enough to all ship sizes and image resolutions.

In the semi-automatic mode, the SUMO software works with a user interface where the user can navigate drop-down menus to select input data, processing parameters and output options. Furthermore, the image and the detections can be viewed at any zoom level and contrast. When run fully automatically, the SUMO software is launched in batch mode with a file in which the processing parameters and input/output options are assigned.

Apart from selection of the input image file and the specification of the output options (output file name, location and format), the processing parameters that can be selected are:

- Nominal false alarm rate ( $P_{FA}$ ),
- Detection threshold adjustment ( $f$ ),
- Equivalent Number of Looks (ENL),
- Land mask buffering.

The nominal false alarm rate  $P_{FA}$  is the expected relative occurrence of sea clutter pixels above the local detection threshold (regardless of frequency, polarization or any other parameter). The CFAR

approach proposes that all sea clutter (non-ship) pixels have values that follow the model PDF that is fitted to the local background. The false alarm rate is chosen to result in an acceptable number of false alarms when analyzing an image. Again considering a reference image size of  $10,000^2$  pixels, and taking into account that SAR images are usually sampled with the Nyquist criterion meaning that there are two samples per resolution element, a typical image has  $(10,000/2)^2$  independent values, each of which may randomly fall above the detection threshold. A false alarm rate of  $P_{FA} = 10^{-7}$  on such an image would give an expected number of  $10^{-7} \cdot (10,000/2)^2 = 2.5$  false alarms in the whole image, which might be acceptable as a compromise between too many false alarms and too many undetected small ships. The  $P_{FA}$  must be chosen based on user requirements. It can be lowered if it is important that false alarms are minimized, even at the cost of missing weak targets; or it can be raised if one wants a more complete result in finding also the weak targets, with the consequence of including more false alarms. However, values of  $P_{FA} \geq 10^{-5}$  that are sometimes quoted would lead to  $\geq 250$  false alarms in a typical contemporary SAR image, which seems excessive.

In SUMO, the three-parameter K distribution is used as the model PDF, in line with its widely-accepted use to describe SAR images of the sea surface [43–45]. The K distribution is the convolution of two gamma distributions, and as a consequence, its PDF has one parameter that determines the mean and two parameters that determine the width. The mean parameter reflects the mean level of radar backscatter from the sea surface (mainly due to local mean wind and incidence angle). The two width parameters reflect the two processes that distribute the pixel values around their average in the radar image. The first process is the speckle noise that is caused by the interference of the coherent radar signals. The level of the speckle noise is determined by the multi-look averaging that was done in the construction of the image and is quantified by the ENL. Therefore, this is a parameter fixed for the whole image, or at least for the whole sub-swath in the case of ScanSAR images. Indeed, the value of the ENL parameter is normally obtained from the metadata, although it can also be estimated from the data [46]. The second process is the intrinsic variation of the radar backscatter level of the sea surface, following physical effects, such as small-scale changes in surface roughness (due to wind) or slope (due to waves). The mean and the intrinsic sea clutter level are calculated from the local pixel mean and standard deviation values (see Section 3 below).

When using the proper ENL value and a suitable  $P_{FA}$  value, it is found that the actual number of false alarms is in most cases significantly higher than what is theoretically expected. The explanation is that the K distribution is not a perfect model of the real data. In particular, the tail of the real data distribution can be longer than that of the model K distribution. To correct for that, an adjustment parameter that raises the detection threshold is used. A suitable value for this parameter was found by experience and can be kept unchanged under many conditions. However, under some conditions, such as high sea state, the presence of icebergs or a very inhomogeneous wind field, the detection threshold will need to be raised above its typical values.

While the K distribution is not perfect, it is still a preferred choice. Many other parametric distributions have been proposed to model radar sea clutter; see the reviews [47,48]. The ones mostly used are Gaussian [12], log-normal and Weibull [44,49,50]. The Gaussian PDF is the simplest, but it does not adequately reflect the asymmetric distribution of the clutter, except maybe at low resolution and high ENL. Although the other two may fit very well in many situations (ranges of wind speed, incidence angle, resolution, etc.), it seems that the K distribution offers an acceptable fit in even broader circumstances (e.g., [51]). Unlike most other parametric distributions, its two components reflect a physical background [43,52] that is, as mentioned, the speckle contribution (which is in fact theoretically of the gamma form for fully developed multi-look speckle) and the intrinsic sea clutter contribution (which in its generality defies any simple parametric model). Recently, the generalized gamma distribution has been found to work well for sea clutter [53,54]. Conceivably, distributions that explicitly include a thermal noise component (e.g., [55,56]) would be well suited for cross-pol images, where the signal level is low. We have not assessed whether any of these are better than the K distribution.



Finally, the land mask buffering parameter sets a seaward extension of the land mask to prevent inaccuracies in the land mask, which do occur, from giving rise to too many false alarms. A typical buffer size can be 100 m.

A particular cause of false alarms in maritime SAR images are azimuth ambiguities [46]. These are image artefacts, repetitions of targets at a much lower level at fixed distances in azimuth [57]. In many cases, the azimuth ambiguities of a target are below the clutter level, but for strong targets and low clutter level, they may be visible and get detected. The offset distance in azimuth where the ambiguities appear is a determinate function of the image parameters, including the Pulse Repetition Frequency (PRF), and is often of the order of 5–10 km. Bright targets on land near the coast may produce azimuth ambiguities on the sea, something that happens often near ports. SUMO checks if there is a brighter target at the known azimuth ambiguity distance at either side of a detected target, and if so, the detection is flagged as a possible false alarm.

A similar effect occurs in range, as range ambiguities; however, in that case, the typical offset distance in range is 150 km, which in most cases puts the possible range ambiguity source outside of the image, precluding any verification. Therefore, SUMO does not check on range ambiguities. However, for a series of repeat-pass images, range ambiguities can often be identified; see Section 4.

### 2.3. Input Data

We can distinguish two types of input data: radar data (the digital radar image plus its metadata) and auxiliary data.

Regrettably, nearly all satellite SARs offer their radar data in a different format. For each new satellite SAR, new ingestion software needs to be implemented. At the time of writing, SUMO can ingest images and their metadata from the satellite SARs listed in Table 1. Note that these systems cover radar frequency bands from L (24 cm) to X (3 cm), and resolutions from 1–100 m.

**Table 1.** Satellite SARs that can be processed by SUMO. In the last column, “I” means that the metadata resides inside the image data file; “I + A” means that the metadata is partly inside the image data file and partly in one or more auxiliary files. In the “Lifetime” column, no end date means that the sensor was still operational at the time of writing. In the “Image data format” column, CEOS refers to the SAR data products format standard of the Committee on Earth Observation Satellites (ceos.org) and HDF5 is Hierarchical Data Format 5 of the HDF Group (hdfgroup.org).

Sensor	Band	Country/Organization	Lifetime	Image Data Format	Metadata Format
ERS-1	C	ESA	1991–2000	CEOS	I + A
ERS-2	C	ESA	1995–2011	CEOS	I + A
RADARSAT-1	C	Canada	1995–2013	binary	I + A
ENVISAT-ASAR	C	ESA	2002–2012	binary	I
RADARSAT-2	C	Canada	2007–	tiff	xml
Cosmo-SkyMed	X	Italy	2007–	HDF5	hdf5
TerraSAR-X (and TanDEM-X)	X	Germany	2007–	MGD: tiff SLC: binary	xml
Sentinel-1	C	EU/ESA	2014–	tiff	xml
ALOS-2 PALSAR	L	Japan	2014–	tiff or CEOS	I + A

SUMO works on images in natural coordinates, range and azimuth. It does not need or indeed want any pre-processing such as map projection, orthorectification, DEM correction or radiometric calibration; nor is any pre-processing with speckle reduction filters deemed desirable; while this may improve detection results for some types of ships and in some clutter backgrounds, we believe that the unfiltered images will provide the most generally optimal results.

As mentioned, SUMO does not use any radar phase information. The kind of phase information that is available from today’s satellite SARs and that is useful for ship detection is polarimetric. While fully polarimetric SAR images have more potential to discriminate ships from the sea clutter background, such images are severely limited in swath and, therefore, less well suited for maritime surveillance. In addition, methods to use fully polarimetric data for ship detection are not yet mature,

and the added value over using dual polarimetry has not been well quantified and may indeed be limited, especially compared to the loss in swath width. In dual polarimetry, the combination of cross-pol (VH) and HH provides the best option: cross-pol gives a very high contrast between ships and sea, but some ships have a low cross-pol RCS, and at far ranges, the cross-pol signal becomes weak; there, the HH channel can still provide a good ship-sea contrast [21,58]. Often, the dual-pol combination cannot be chosen, and then, cross-pol plus VV is also suitable, but the ship-sea contrast in VV is much less than in HH, especially for higher wind speeds. In the future, hybrid or compact polarization [59,60] may provide suitable data for dual channel ship detection [61]. Although not used today by SUMO, the phase difference between the dual polarization channels might help in the detection process. On the other hand, single-to-noise would be a limiting factor just like it is now for the cross-pol amplitude, and some techniques that exploit polarimetric phase involve a spatial convolution (e.g., [23]), which would result in the loss of small targets.

SUMO can ingest both Multi-look Ground-range Detected (MGD) images, which are in ground range projection and where the pixel values represent the radar backscatter amplitude, and Single-Look Complex (SLC) images, which are in slant range projection and where the pixels have complex (real and imaginary) values. In the SLC case, however, the complex pixel values are as a first step converted to amplitudes, and from there, the processing is no different from the MGD processing (strictly speaking, to ensure adequate sampling, the complex image ought to be first interpolated before taking the amplitude, but for efficiency, this is not done). There is no geometrical transformation of the SLC image from slant to ground range and no multi-looking, but in the end, the derived attributes (location, size, heading, etc.) are corrected for slant range distortions.

As described in the previous section under Point 6, multi-polarization images can be processed, treating each polarization channel separately, disregarding the phases and combining the separate results from the channels after pixel-based detection with a union (“logical or”) operation. The motivation for this approach is that individual scatterers on the ship can be different ones in HH, VV or cross-pol, and in this way, all are equally combined to better fill in the outline of the ship. In the literature, other approaches to incoherently combine multiple polarization channels have been proposed: Cross-correlation [62], multiplication [63,64], principle component analysis [65] or summation after normalization [66].

Concerning auxiliary data, SUMO ingests the following:

- Coastline vector data for the land mask,
- Ice edge vector data,
- Any other vector data for overlay display purpose,
- Previous SUMO output data files.

Suitable coastline vector data are the open source datasets GSHHG/GSHHS (Global Self-consistent, Hierarchical, High-resolution Geography/-Shoreline) [67] and OpenStreetMap [68]. For faster ingestion, these global datasets have been cut up into geographic tiles; only the tiles that cover the radar image under analysis need to be ingested. As a further measure to increase speed, a set of pre-buffered land masks is pre-produced for a few often used buffer sizes like 50 m, 100 m and 250 m. As the buffered land masks need fewer vertices, this does not lead to much extra data and speeds up the land masking.

The ice edge data are used in a similar way as the land mask, to mask out areas of sea ice. These data, however, change on a daily basis. Suitable daily ice edge data that are automatically ingested on-line by SUMO are provided by NOAA/the U.S. Navy [69].

Other vector data (points, lines or polygons) that can be overlaid to aid in the human interpretation of the detections can, e.g., include ship positions from other sources or elements from nautical maps. Overlaying detections from previous images can help to understand whether a target has a continued presence.

## 2.4. Output Data

One image file gives rise to one set of results, independent of any other images analyzed. The results are output in an xml-formatted file that contains three sections: the first with some image metadata, in particular sensor, image mode, corner coordinates and start/stop times; the second with the parameter values used in the ship detection; and the third with the list of detections and their attributes.

The attributes for each detection are the following:

- Row and pixel number;
- Geographic longitude and latitude in WGS84;
- Time of detection (can be many seconds off from the image starting time);
- Length and width in meter;
- Heading in degrees w.r.t. the range direction and w.r.t. north;
- Number of detected pixels in the target signature;
- Maximum pixel value and detection significance;
- Radar Cross-Section (RCS);
- Reliability figure.

Location information refers to the center of the target signature. Heading has an ambiguity of 180°, because bow and stern are not distinguished. The detection significance is defined as (NB, the background estimation will be discussed in Section 3):

$$\frac{(\text{Maximum pixel value of the target} - \text{mean of background pixel values})}{\text{standard deviation of background pixel values}} \quad (1)$$

As the PDF of the clutter has a much longer tail than a Gaussian, and the more so for images with lower ENL values, the significance must be relatively high for a statistically-reliable detection. The RCS is an integrated property of the target and requires the radiometric calibration, which is obtained from the metadata. Finally, the reliability figure has four possible values: very likely to be a false alarm, probably a false alarm, probably a true ship, very likely to be a true ship. These can be seen as the likelihood of being a true ship at 15%, 40%, 70% and 95% (see also Section 3.8 below). While the precise likelihood boundaries are arbitrary, these four classes have been found useful in deciding which detections to include in the final result presented to a user, depending on the application. The stochastic nature of the pixel values in a SAR image together with the occurrence of physical effects that may mimic a ship signature lead to intrinsically uncertain results, hence the need for this attribute.

Furthermore, a simple file for Google Earth viewing in KML (Keyhole Markup Language) format is produced.

## 3. Detailed Aspects of the Detection Algorithm

### 3.1. Image Edges

Many systems provide an image file with invalid pixels near the edges and corners. The pixel values in these areas can be zero or they can slowly ramp up between zero and the well-defined values. This gives rise to errors in the estimation of the background statistics. SUMO removes low-value edge areas by applying adaptive thresholding to each line in a sparse sub-sampling of image lines along the azimuth and range directions. The aim of the thresholding is to find, for each line in the sample, the location of the boundary between valid and non-valid pixels. The boundary value of one line is then also applied to neighboring (unsampled) lines. This method is based on the observation that, at local level, the values of non-valid pixels are always smaller than the values of the valid pixels. The method does not rely on any image metadata parameter (which are generally lacking for this aspect) and has shown a strong performance, both in removing non-valid pixels, as well as in keeping valid pixels.



### 3.2. Background Estimation

The level of sea clutter varies across the image due to variations in incidence angle, wind and other meteorological and oceanographic effects. It is not meaningful to set an overall threshold, and indeed, the CFAR approach requires the estimation of the local clutter level. The smaller the area for the background clutter estimation, the more adapted it is to variations, but the higher the estimation error in the stochastic quantities mean and standard deviation. As a compromise between accuracy and speed, SUMO splits the image into square tiles for which the statistics are computed independently. For a similar error in the estimate of the standard deviation, more independent samples from the distribution are needed than for the estimate of the mean. To deal with that, SUMO splits each tile into four sub-tiles and computes the mean ( $M$ ) in each sub-tile, normalizes each sub-tile by dividing by its mean and then computes a single standard deviation ( $S$ ) from the normalized pixel values from all four sub-tiles together. In this way, a measure for the clutter PDF width is found for each tile, while the mean is taken out by the normalization.

Of course, the use of such tiles is only a rough way to deal with the background variations. One source of background variation is the drop-off of sea surface backscatter with range and incidence angle. This behavior is to some extent deterministic, so it could be considered to use a model to account for this. However, such a model would still be non-trivial and only approximate, because the function of radar backscatter on incidence angle depends on a series of imaging parameters, many of which including the main one, wind velocity, are not exactly known and vary locally. Therefore, no attempt is made to include a range-dependent model component in the background estimation.

Experimentation with real and simulated data has shown that a tile size of  $200 \times 200$  pixels gives optimal results. A smaller tile size results in variations of the estimated parameters that are dominated by random noise. For a larger tile size, intrinsic background variations start to dominate. Ideally, the tile size should be adapted to the spatial scales of the background clutter variations, but this is not done (also to save computing time). The chosen tile size implies the use of 2500 independent samples for the mean (sub-tile of  $100 \times 100$  pixels that contains  $50 \times 50 = 2500$  independent values following the Nyquist sampling) and 10,000 independent samples for the standard deviation (from the four sub-tiles together). Only every second pixel in range and in azimuth is used, as neighboring pixels are not independent. In the case where part of the tile is covered by the land mask, fewer samples are used, but should too few samples remain in a tile after land masking, clutter PDF values from a neighboring tile are used.

### 3.3. Threshold Setting

This section describes the conceptual approach for the calculation of the CFAR detection threshold, but the actual implementation is further modified in the next section.

The three parameters of the K distribution are usually denoted as  $\mu$  for the PDF mean and  $L$  and  $\nu$  for the PDF width measures.  $L$  and  $\nu$  are symmetric,  $L$  denoting the (known) ENL leading to the speckle level and  $\nu$  measuring the inhomogeneity of the intrinsic (speckle-free) sea surface backscatter. A (sea) surface with a constant backscatter level is represented by  $\nu = \infty$ , which leads to a K distribution limit equal to a gamma distribution, while  $\nu = 1$  characterizes the most inhomogeneous sea surface backscatter.

According to the model for speckle and sea clutter, it is the radar backscatter power values that are characterized by the K distribution. The pixel values in the usual radar image product however are amplitudes, being the square root of the powers, save for a calibration factor. SUMO works only with amplitude values. Given that power values are K distributed, the PDF of the corresponding amplitude values is also known. We will write  $\mu$  and  $\sigma^2$  to denote the intrinsic (unknown) mean and variance of the amplitudes of K distributed clutter, while for their estimates based on pixel (amplitude) values  $x$ , we will use  $M = \langle x \rangle$  and  $S^2 = \langle (x - M)^2 \rangle$ . From  $\mu$  (estimated by  $M$ ),  $\sigma$  (estimated by  $S$ ) and  $L$  (provided),  $\nu$  can be derived from theoretical relations (e.g., [70]).

With  $\mu$ ,  $L$  and  $\nu$  known, the CFAR (amplitude) threshold  $\vartheta_{FA}$  can be calculated from the integral equation:

$$\int_{\vartheta_{FA}}^{\infty} f_K(a) da = P_{FA}, \quad (2)$$

with  $f_K(a)$  the amplitude PDF of  $K$  distributed sea clutter. This can be numerically solved for  $\vartheta_{FA}$  given a choice of  $P_{FA}$ . As  $P_{FA}$  is very small, special care has to be taken for an accurate numerical implementation. In case of very low or very high values for  $L$  and  $\nu$ , limiting approximations for the distribution need to be used.

Solving the integral equation takes some computing time. For that reason, the equation is pre-solved for a wide set of  $(L, \nu)$  values, and the results are tabulated. During actual operation, this lookup table is used with linear interpolation.

The “detection threshold adjustment” parameter  $f$  that is used to compensate for the difference between the real clutter PDF and the fitted  $K$  distribution is implemented as:

$$\vartheta'_{FA} = (\vartheta_{FA} - M) \cdot f + M, \quad (3)$$

or in terms of values normalized by the background clutter mean:

$$\vartheta'_{N,FA} = (\vartheta_{N,FA} - 1) \cdot f + 1. \quad (4)$$

The  $\vartheta'_{N,FA}$  is the CFAR detection threshold that is finally applied. Experience based on a very large number of images drawn from all systems mentioned in Table 1 has shown that for co-pol channels (HH, VV), a value of  $f = 1.5$  works well under many circumstances and for cross-pol channels (HV, VH), a value of  $f = 1.2$ . In the case of images with strong sea clutter (e.g., many breaking waves), these values need to be raised to prevent excessive false alarms, maybe to 1.8 or 2.0 for co-pol. Indeed, it can be argued that breaking large waves cause real individual backscatter events that are not modelled by the  $K$  distribution.

### 3.4. Avoiding Contamination by Non-Sea Pixels in the Background Estimation

The tile may contain ships or unmasked slivers of coast, which have high pixel values that are not part of the sea clutter PDF. These are removed by an iterative procedure. First, the normalization by the mean is performed, and the standard deviation of the normalized pixel values is calculated, as described above, using all pixel values in the tile. Next, a clipping threshold,  $\vartheta_{CL}$ , is calculated similar to the false alarm threshold,  $\vartheta_{FA}$  above, but using instead of  $P_{FA}$  a relatively high  $P_{CL}$  of 0.05, i.e.,

$$\int_{\vartheta_{CL}}^{\infty} f_K(a) da = P_{CL}. \quad (5)$$

Then, all pixels above the clipping threshold are removed, meaning that nominally, the 5% highest pixels are removed. With the remaining pixels, the procedure is repeated. This is done several times, until the results do not change much anymore. Usually, three iterations are enough; the current SUMO implementation uses a fixed number of two iterations.

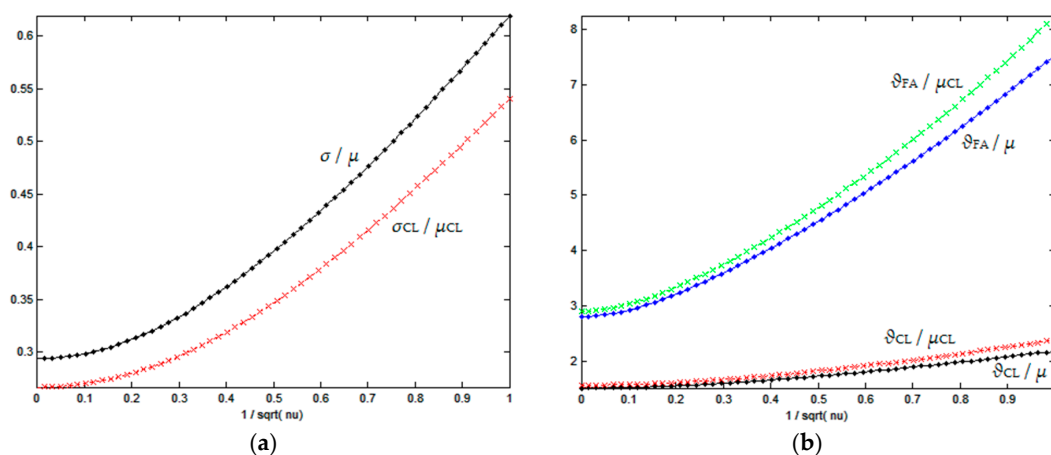
Calculating the mean and standard deviation using only the 95% lowest values subset of the PDF obviously gives a bias, leading to too low values of the estimated parameters. However, the bias is known, and it can be calculated by numerically integrating the PDF. Therefore, we numerically compute for a grid of  $\nu$  and  $L$  values:

$$\mu_{CL} = \int_0^{\vartheta_{CL}} a f_K(a) da / \int_0^{\vartheta_{CL}} f_K(a) da \quad (6)$$


$$\sigma_{CL}^2 = \int_0^{\vartheta_{CL}} (a - \mu_{CL})^2 f_K(a) da / \int_0^{\vartheta_{CL}} f_K(a) da \quad (7)$$

as well as the  $\vartheta_{FA}$  and  $\vartheta_{CL}$  defined above. With this, for a certain  $L$ , we can relate a given  $\sigma_{CL}/\mu_{CL}$  to a  $\vartheta_{FA}/\mu_{CL}$  and  $\vartheta_{CL}/\mu_{CL}$ , as these all vary with  $\nu$ . Writing  $x_N = x/M_{CL}$  and  $\vartheta_N = \vartheta/M_{CL}$  for normalized amplitude values and using  $M_{CL} = \langle (x < \vartheta_{CL}) \rangle$  as the estimator for  $\mu_{CL}$  and  $S_{N,CL} = \sqrt{\langle (x_N < \vartheta_{N,CL}) - 1 \rangle^2}$  as the estimator for  $\sigma_{CL}/\mu_{CL}$ , a value of  $\vartheta_{CL}/\mu_{CL}$  can be looked up to use as clipping threshold  $\vartheta_{N,CL}$  on the normalized pixel values in the iterative procedure and, finally, a value of  $\vartheta_{FA}/\mu_{CL}$  to use as detection threshold  $\vartheta_{N,FA}$ .

In the numerical calculations, the grid of  $\nu$  values is sampled with equidistance spacing in  $1/\sqrt{\nu}$ . This gives a more homogeneous sampling of the domain of the variables of interest than an equidistant sampling in  $\nu$ . The  $L$  values used are those that correspond to the ENL values of existing image products. Figure 1 gives an example of some intermediate numerical calculation results for  $L = 3$  as a function of  $\nu$ . Figure 2 shows the lookup tables for five different values of  $L$ .



**Figure 1.** Numerical calculations of variables as a function of  $1/\sqrt{\nu}$ . (a)  $\sigma/\mu$  (black/upper) and  $\sigma_{CL}/\mu_{CL}$  (red/lower). (b)  $\vartheta_{FA}/\mu_{CL}$  (green/highest),  $\vartheta_{FA}/\mu$  (blue/second),  $\vartheta_{CL}/\mu_{CL}$  (red/third),  $\vartheta_{CL}/\mu$  (black/lowest).



**Figure 2.** The direct lookup tables from the clipped standard deviation divided by the clipped mean ( $\sigma_{CL}/\mu_{CL}$ , which is calculated from the data as  $S_{N,CL}$ ), to (a) the clipping threshold divided by the clipped mean  $\vartheta_{CL}/\mu_{CL}$  and (b) the detection threshold divided by the clipped mean  $\vartheta_{FA}/\mu_{CL}$ , which can both be applied directly to the data as  $\vartheta_{N,CL}$  and  $\vartheta_{N,FA}$ . Five plots are included, for values of  $L = 1, 2, 3, 10, \infty$ , from right to left, in black, blue, cyan, green and red. Each of these has only a limited range of existence in  $\sigma_{CL}/\mu_{CL}$ .

### 3.5. Clustering

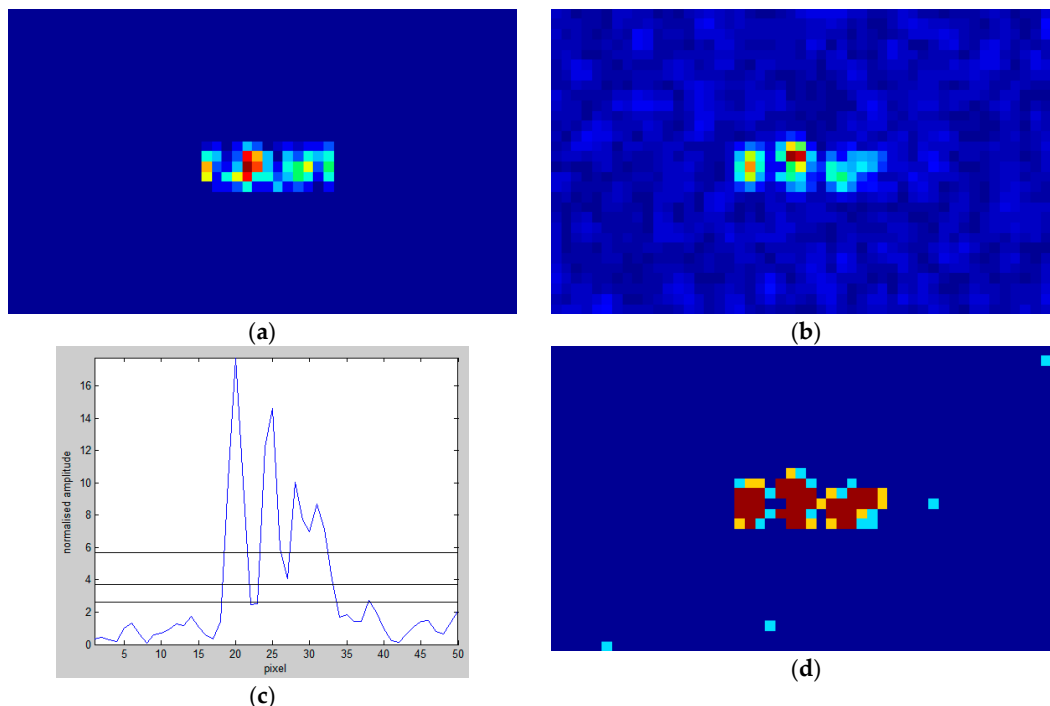
Individually-detected pixels of the same ship need to be grouped together (as mentioned in Section 2.2, for a multi-polarization image, a pixel is detected if it is above the detection threshold in any channel). When a ship extends over many pixels (ship size > resolution), not all pixels within the ship's geometrical outline may be bright, and in particular, not above the detection threshold, as the brightness of a pixel depends on the presence of particular reflecting structures within the pixel and their interference (fading). Therefore, grouping of neighboring detected pixels would lead to fractured objects. At the same time, neighboring pixels are not independent: as mentioned before, the images are sampled at two samples per resolution element (Nyquist), leading to increased brightness around any individual pixel with a high intensity. In order to group all detected pixels from one object into a single cluster that is representative of the shape of the object, a new threshold called the clustering threshold is defined, lower than the detection threshold, and it is this threshold that governs the grouping of neighboring pixels. The clustering threshold is set to the mean plus three standard deviations of the clutter background, where now, the background refers to a new  $200 \times 200$  pixel tile that is centered on the detection. The same iterative clipping technique as before to avoid contamination from high-valued pixels is repeated on the new tile. The clusters are created by cluster growing starting from the highest detected pixel that is not yet included in a cluster, repeating until all detected pixels are in a cluster. In that way, all detected pixels will be part of a cluster of surrounding pixels above the clustering threshold, while pixels that are above the clustering threshold, but below the detection threshold and not connected to a detected pixel will not be included in any target cluster. As neighbors, all eight pixels up, down, left, right and diagonal are used (eight-connectivity). Cluster growing is stopped if the physical size of the cluster (in  $\text{m}^2$ ) becomes unrealistically large for a ship; in such a case, probably an unmasked part of land or an extended sea surface feature was detected.

The cluster threshold is relatively low in order to bridge gaps between bright parts of the same object. It tends to extend beyond the target outline, including nearby sea clutter pixels that are accidentally somewhat high. In order to obtain a better target outline, a third threshold is defined, the signature threshold. Its level is set at the mean plus five standard deviations of the clutter background, putting it in-between the clustering threshold and the detection threshold. All pixels of the target cluster above the signature threshold are selected as being part of the target signature, which is taken as the final estimate of the target. Note that the target signature may contain unconnected pixels. Furthermore, note that the same target signature applies to all of the polarization bands in the image. An example of a clustering is given in Figure 3, and the thresholds are summarized in Table 2.

The levels of the clustering and signature thresholds (three and five sigma) were determined by trial, as was the use of eight-connectivity to define a neighbor as opposed to four-connectivity (only up, down, left and right). These choices determine to what extent neighboring bright areas will be merged into one object and are to some extent arbitrary. However, this is closely related to the nature of the radar signatures of ships. It is fundamentally impossible to distinguish a single long ship that in a radar image only shows strong scattering from its bow and stern, but not from its midsection, from two separate smaller ships.

**Table 2.** Overview of the thresholds used, from high to low.

Threshold	Level	Contains
Detection	$\vartheta_{\text{FA}}$ modified by adjustment factor $f$ , calculated from original tile	Detected pixels
Signature	mean + 5 SD of clutter background calculated from tile centered on target	Target (ship) signature
Clustering	mean + 3 SD of clutter background calculated from tile centered on target	Target (ship) cluster



**Figure 3.** Illustration of the clustering. Only amplitude values are visualized, although (a,b) are from complex data. (a) Model backscatter distribution of a ship. Each pixel represents a point scatterer, the color from blue to red signifying increasing amplitude, with random phase (not visible). (b) Simulated SAR image of the ship on a sea clutter background with  $\nu = 3$ ,  $L = 1$ , SNR (target maximum/sea clutter mean) = 71 and using a Kaiser taper function [71] in the spectral domain in both directions with parameter 6.4 and support 90%. The backscatter intensity distribution has changed from the model (a) due to fading between neighboring ship pixels. (c) Horizontal cross-cut through the center of simulated image with the three thresholds superimposed, from low to high: clustering threshold (clutter mean + 3 sigma), signature threshold (clutter mean + 5 sigma), CFAR detection threshold. (d) Clustering result. Red: pixels above the CFAR detection threshold. Orange: ship signature. Cyan in contact with red: ship cluster.

### 3.6. Estimation of Size, Heading and RCS

Ideally, and indeed often, the target signature of a ship as derived above looks like a ship's outline being elongated (ellipsoidal/rectangular) in shape. The size and heading of the ship are estimated by treating the pixels of the signature as points on a Cartesian (x, y) grid and fitting a least-squares line to this point cluster. The orientation of the fitted line w.r.t. the range direction is taken as the ship's heading w.r.t. range (with  $180^\circ$  ambiguity). The distance between the extreme pixels along the fitted line is taken as the length and the distance perpendicular as the width.

Recently, an improved method for ship size estimation was developed [72]; however, this is not yet implemented in SUMO.

Ships of a size smaller than or similar to the image resolution, inasmuch as detected, will show up as a single point or a small roundish signature. Derived length and width will be similar, and the derived heading will be unreliable.

The RCS of a target measures the total amount of radar energy it returns. It is calculated by the sum of the squares of the pixel values (i.e., radar power values) over all pixels in the target cluster and applying the calibration constant provided in the metadata. Because of fading, especially for targets that are small with respect to the resolution, the RCS value of a single observation depends very strongly on the measurement geometry [73] and is therefore only weakly representative of physical target RCS (which would need to be described as a function of aspect and incidence angles).



### 3.7. Geographic Location

The image metadata provide the means to transform a (pixel, row) location in the image to a geographic (longitude, latitude). Usually, this is done by interpolation in a grid of reference points or by polynomial transformation. Although very accurate geographic positions are generally not needed for the ships (if they are moving, their position will, of course, quickly change), they are needed to overlay high-resolution land masks and to recognize stationary targets from one image to the next. For Sentinel-1, it was found that using the reference grid did not always lead to satisfactory results [74], so a more accurate method was developed that reaches the required accuracy.

If the orbit state vectors included in the image metadata are accurate enough (as is the case of the Sentinel-1 images delivered with standard timeliness category), these vectors can be used to give a very precise geolocation of the image pixels. The geolocation method solves the range Doppler equations to convert image coordinates into geographic coordinates, or vice versa. In the case of forward geocoding (from image coordinates to geographic coordinates, e.g., to know the geographical location of a detected ship), the azimuth coordinate of the Pixel of Interest (PoI) is first translated to zero Doppler azimuth time. This is easily accomplished since in modern sensors, the image azimuth lines are annotated with their zero Doppler azimuth time in the metadata. The satellite position (S) and velocity at that azimuth time are then calculated by interpolation of the orbit state vectors. The PoI's geographic coordinates will be the only coordinates that meet all of the following constraints:

1. The PoI is located on the plane that is perpendicular to the orbit trajectory at point S. This is because only the points on the plane perpendicular to the satellite's trajectory will have a zero Doppler shift.
2. The PoI is located at a specific distance from S. This distance can be calculated from the slant range distance to the near range of the image (given in the metadata) and the PoI's image range coordinate.
3. The PoI is on the Earth's surface. Since we are only interested in points on the sea (i.e., ships or coastal locations), we use a geoid model (EGM96—Earth Gravitational Model 1996) as a simplified representation of the Earth's surface.

The reverse geocoding case (from geographic coordinates to image coordinates, e.g., to know the image location of the land mask shapefile) is based on similar principles as the forward one. In this case, the satellite position (S) that will give zero Doppler at the PoI can be calculated as the position along the interpolated orbit trajectory that is geometrically closest to the PoI. Once S is found, the PoI's zero Doppler azimuth time is also obtained, which can then be used to derive the PoI's azimuth coordinate on the image. The distance from S to the PoI is used to derive the PoI's range coordinate on the image.

SAR images have the characteristic of showing the radar echo from a moving object at a location that is displaced in the azimuth direction by an amount that is proportional to the radial velocity of the object w.r.t. the radar. For ships moving on the sea surface, the azimuth displacement  $D_{az}$  is given by:

$$D_{az} = (R_{slant}/V_{sat}) \cdot \sin(\vartheta) \cdot \cos(\phi) \cdot V_{ship} , \quad (8)$$

where  $R_{slant}$  is the slant range distance between the satellite and target,  $V_{sat}$  the satellite platform speed in orbit,  $\vartheta$  the incidence angle,  $\phi$  the vessel course with respect to the range direction and  $V_{ship}$  the speed of the ship. This means that moving ships will be found not in their actual location, but slightly offset. The geographic coordinates output by SUMO give the apparent (offset) position, as is seen in the SAR image. It is not possible to correct to the real position, as the ship velocity is not known (SUMO does not attempt to estimate a target's speed). Considering some typical values ( $R_{slant} = 950$  km,  $V_{sat} = 7.5$  km/s,  $\cos(\phi) = 1$ ,  $\vartheta = 45^\circ$ ), the offset for a ship moving with  $V_{ship} = 15$  knots is 700 m. This effect is indeed observed in the offset of a ship's location from its wake (if the wake is visible) [75] and is important to take into account when correlating with ship reporting data (see Section 4).

### 3.8. False Alarm Discrimination and Reliability

The check for azimuth ambiguities was described above in Section 2.2. The expected distance for azimuth ambiguities is [57]:

$$D_{\text{azamb}} = m \cdot \lambda \cdot \text{PRF} \cdot R_{\text{slant}} / (2 \cdot V_{\text{sat}} (1 - \cos(\psi) / \text{RPD})) , \quad (9)$$

where  $\lambda$  is the radar wavelength,  $m$  the order of the ambiguity and the factor after  $V_{\text{sat}}$  containing the orbit inclination  $\psi$  and the number of revolutions per day (RPD) is a small correction for Earth's rotation. The present implementation of SUMO only checks for azimuth ambiguities at  $m = \pm 1, 2$ , although in fact, higher order ambiguities are sometimes seen.

Sentinel-1 images display an additional azimuth ambiguity-like feature [74,76]. It is also found to be at a determinate distance, and it is similarly treated.

If a detection is deemed to be an azimuth ambiguity, it is flagged as such, and its reliability is put to the lowest level. In addition, the reliability figure is reduced from its maximum possible value if any of the following conditions occur:

- Improbable (too high) length or width, taking into account maximum possible ship sizes;
- Improbable (too low) length-to-width ratio, but only if the target is well resolved;
- Low significance.

While the decision boundaries are based on physical considerations, their exact values and how much the reliability is lowered are somewhat arbitrary, reflecting the stochastic nature of the radar imaging.

In the final output, the detections with the lowest reliability level can be considered false alarms. Still, it is useful to retain them, because independent information (such as a reported ship position in the same location; see below) may still vindicate such a detection.

### 3.9. Parameters

SUMO aims to follow where possible a rigorous approach, based on the inherent characteristics of the SAR imaging. Still, many aspects of the target and clutter SAR signatures are intractable, and the models are only rough. This results in the use of a number of parameters for which appropriate values need to be chosen. This is acceptable if the parameter values can remain constant over a wide range of conditions, and experience has shown that this is indeed the case. Table 3 gives an overview of all parameters used in SUMO as discussed in the text. The quoted values are those that have been determined to work best from experience on a large set of images. Only the first three parameters are real user parameters that can be adjusted during operations, following the indications in the 'Remarks' column. The other parameters do not need adjustment. Their values can remain the same for all SAR sensors tested (Table 1), radar bands (X, C, L), incidence angles, polarizations (with the distinction in  $f$  as noted), wind speeds, sea states, geographic areas and ship types. The only exception is for very high sea state or in situations where many false alarm sources are present in the image, where  $f$  needs to be raised (these situations are further specified in Section 6). In the case where large targets are imaged at high resolution (e.g., a 400-m ship at 1-m resolution), the tile size for background estimation would become too small. However, such a combination should be avoided in the first place, because it represents the wrong choice for the task of ship detection; it is, instead, better suited for ship classification, which SUMO does not incorporate.

**Table 3.** Parameters used in Search for Unidentified Maritime Objects (SUMO).

Parameter	Symbol	Value	Remarks
Nominal false alarm rate	$P_{FA}$	$10^{-7}$	Can be raised in case being more complete in detecting weak targets outweighs the increase in false alarms.
Detection threshold adjustment	$f$	1.5 (co-pol) 1.2 (cross-pol)	Increase if sea state is very rough or if many false alarm sources are present.
Land mask buffer size		100 m	Increase if coastline is not well matched by land mask.
Equivalent Number of Looks	ENL	From metadata	Specify if not present in metadata.
Tile size for background estimation		$200 \times 200$ pixels	Balance between random noise and scale of clutter inhomogeneities.
Tile size for target clustering background		$200 \times 200$ pixels	As above.
Clipping rate to remove high pixel values from background estimation	$P_{CL}$	0.05	Balance between excluding potential target pixels and including high values of clutter PDF.
Number of iterations to remove high pixel values		2	Balance between accuracy and speed.
Clustering threshold		3 sigma	To link parts of the same ship signature.
Signature threshold		5 sigma	To exclude surrounding clutter pixels from ship signature.
Reliability brackets		15%, 40%, 70%, 95%	Indicative numerical values to divide the detection reliability into 4 classes from low to high.

#### 4. Embedding

The ship detection result of a certain satellite radar image that is output from SUMO is itself input for further processing. To make the context in which SUMO operates clearer, some elements of this are mentioned. At JRC, these functionalities are performed in the Blue Hub, which is the in-house maritime surveillance R&D platform [77].

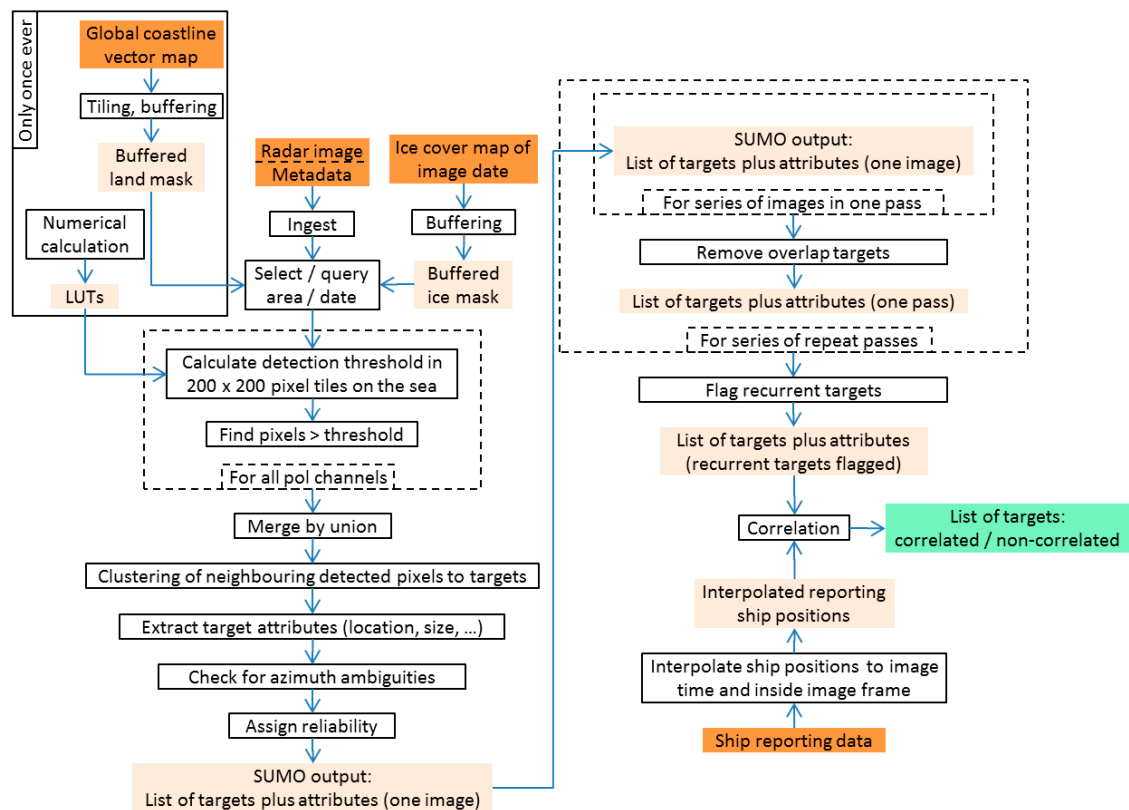
A single satellite image is typically part of a longer imaged swath. Single images are extracted from the entire swath with a certain overlap area in the azimuth direction. If such a series of images is analyzed by SUMO, ships in the overlap areas will be detected twice. Post-processing needs to correct for this.

Many ships use self-reporting systems, such as AIS (Automatic Identification System), LRIT (Long-Range Identification and Tracking) or VMS (for fishing ships). Ship detections from non-cooperative sensors, such as satellite imaging radar, should be fused with the maritime picture obtained from the ship reporting systems, which allows the indication of which of the radar-detected ships are self-reporting and which are not. The fusion process involves interpolation/extrapolation of the reported ship positions to the time of the satellite imaging and correlating interpolated reported positions with detected ship positions, taking into account uncertainties in the interpolation and positioning. Furthermore, the detected ship's azimuth displacement in the SAR image (see Section 3.7) must be accounted for, which is now possible because the ship speed is known from its self-reporting data.

Some prominent false alarm sources include small islands or reefs that are not in the land mask, off-shore constructions, unrecognized azimuth ambiguities and range ambiguities from strong scatterers on land, such as cities or certain mountain slopes [37]. Such false alarms tend to recur in the same position every time the same scene is imaged again. The fixed orbit pattern of a satellite leads to repeated imaging of the same scene in exactly the same geometry, as, e.g., exploited by interferometry. Especially for Sentinel-1, which acquires imagery over the same designated areas in each overpass, many repeat-pass images are available. This makes it possible to flag recurrent targets and indicate them as false alarms in the post-processing [78].


## 5. Flow Chart and End-To-End Example

Figure 4 recapitulates the ship detection algorithm in a flow chart.



**Figure 4.** Flow chart of the SUMO algorithm and its environment. Outlined boxes are processes; text on a color background without an outline signifies data. A dashed box contains a process that is to be repeated as mentioned. The four dark orange colored boxes are the external data that go in. The green box on the right is the final result that comes out. The left half is SUMO proper, which works on a single SAR image. The box at the top left contains a section that needs to be executed only once, whereas the rest is run for each image again. Land mask and ice mask are of course only needed if there is land/ice in the image. The right half depicts the use of multiple images in a maritime surveillance campaign and the combination with ship reporting data.

The rest of this section presents an end-to-end example of an analysis by SUMO of one satellite SAR image. It also illustrates how SUMO's output is combined with ship self-reporting data (AIS) on the Blue Hub platform, which is then used to show a non-exhaustive and mostly qualitative evaluation of SUMO's performance. The image under study is a dual-pol (VV + VH) Sentinel-1 acquisition over Spain's Mediterranean coast, taken on 4 July 2016. The image mode is IW GRDH (Interferometric Wide, Ground Range Detected High resolution [79]), which implies a resolution of  $20 \text{ m} \times 22 \text{ m}$  (ground range  $\times$  azimuth) at an ENL of 4.4. Figure 5 displays the VV polarization of the image, with the OpenStreetMap coastline, buffered by 250 m and used by SUMO for land masking, overlaid in green.



**Figure 5.** VV polarization of the Sentinel-1 image used in this example. The area is the Mediterranean coast of Spain between the cities of Valencia (**bottom left**) and Tarragona (**top right**). Image size is 250 km  $\times$  167 km. A coastline from OpenStreetMap buffered by 250 m is overlaid in green. The green lines on land reflect the tiled storage structure of the pre-buffered coastline. The Sentinel-1 image ID is: S1A\_IW\_GRDH\_1SDV\_20160704T175421\_20160704T175446\_012002\_012865\_965A.SAFE.

Except for the increased land mask buffer size, all parameter values used are those mentioned in Table 3. The automatic SUMO analysis gives 230 detections in the image. Out of these, 82 are deemed azimuth ambiguities by SUMO and automatically assigned the lowest reliability level. Such a large number is typical for coastal areas that often show many azimuth ambiguities from port and urban constructions. The remaining 148 detections are considered for comparison with the reporting ships. There are 32 interpolated reported ship positions (from AIS) inside the image and outside the land mask. A number of reporting ships are docked at the various ports in the image, but those ships will not be detected by SUMO since the ports fall inside the land mask. These reporting ships are therefore excluded from the evaluation. The fusion process between the 148 SUMO detected ships and the 32 interpolated AIS ship positions returns 25 correlations. This number represents 17% of all the detections with higher than the lowest reliability level and 78% of all the interpolated reported positions outside the land mask. Additionally, 123 detections and seven interpolated reported positions are left uncorrelated. Table 4 lists these numbers, and Figure 6 plots them on a map.

After visual inspection of the automatic results, 72 of the 123 detections that are not correlated to reported positions are believed to be not due to ships. Most of them, 34, result from aquaculture mussel farming sites close to the coast, as confirmed by checking on Google Earth. Even if they are not ships, detecting such structures can be of interest for some maritime surveillance users. Work has been performed on recognizing fish cages in SUMO with Radarsat-1 images [30], but it has not been validated for a wide set of sensors, and fish cages look a bit different than mussel farm structures. Range ambiguities are thought to be responsible for 19 detections, identified thanks to the analysis of repeat-pass images, which highlights their recurrent nature. Ten detections are thought to be from buoys, recognized as such because they map out harbor approaches, and confirmed on Google Earth. Unmasked coastal infrastructure (piers) leads to four detections. Based on their shape and brightness, two detections are attributed to oil platforms. Two detections are interpreted to be azimuth ambiguities from ships that are outside the image. Finally, one detection is on a bright ship wake (occurring in



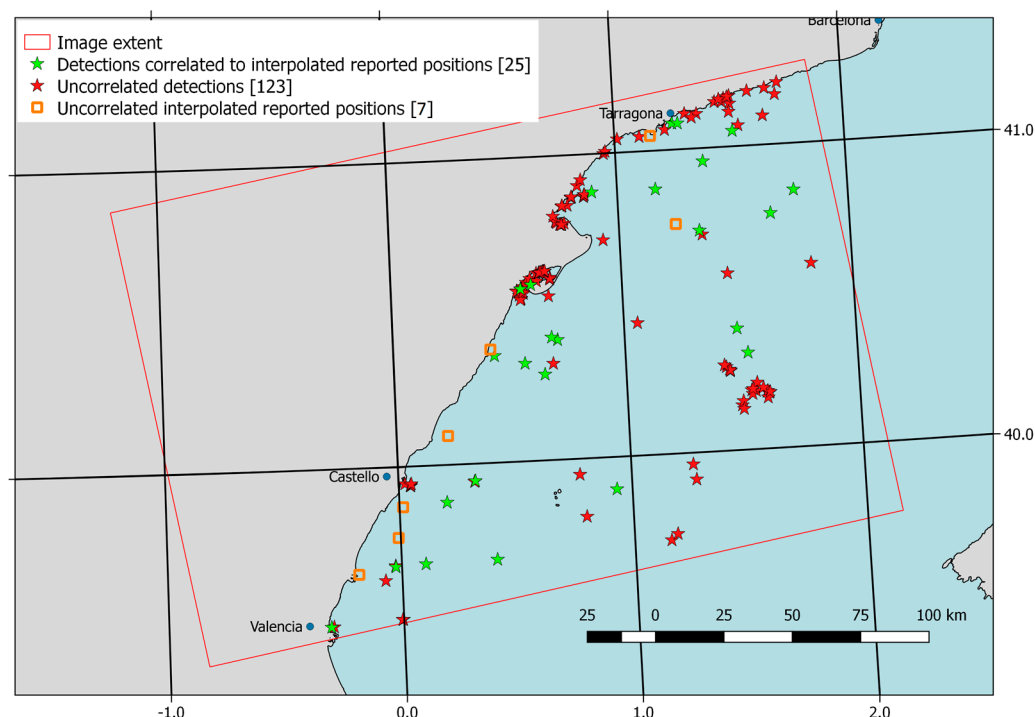
addition to the detection of the ship itself). Figures 7 and 8 shows some examples. Inasmuch as they stay in the same place, most of these, viz. the aquaculture sites, range ambiguities, buoys, piers and platforms, to a total of 69 out of the 72, can be recognized as fixed structures in repeat-pass imagery.

For the remaining 51 of the 123 uncorrelated detections, the visual inspection gives no compelling reason to believe they are not a boat or a ship. Of course, some of them may still be false alarms (recurrent target analysis might weed out a few). If we accept these numbers, the fraction of non-reporting ships out of the detected ships in the area is  $51/(51 + 25) = 67\%$ . This is an acceptable number, because only the larger ships report on AIS, while we may expect many small boats near the coast.


With regard to the seven interpolated reported positions not correlated to detections, three of them are small boats according to their AIS (12, 13 and 24 m in length), which are apparently below the limit of detectability for SUMO (at the  $P_{FA}$  used). More than that, the visual inspection confirms that they are below the limit of “operational detectability” of the Sentinel-1 image: even if some of them do show a very weak signature, it is always at the level of the clutter, so such ships cannot be found in the image without knowing where they are in the first place. The remaining four uncorrelated AIS ships have very long interpolation times, of 1.1, 3.3, 4.9 and 5.4 h, leading to high uncertainties in their interpolated positions.

**Table 4.** Summary of ship detection and correlation results for the example image. Only interpolated reported ship positions outside the land mask are taken into account.


	Number	Fraction	Fraction
Ship detections	230		
Detections with the lowest reliability level	82		
Detections with higher reliability level	148	100%	
Interpolated reported positions	32		100%
Detections correlated to interpolated reported positions	25	17%	78%
Uncorrelated detections	123	83%	
Uncorrelated interpolated reported positions	7		22%



**Figure 6.** Map displaying the SUMO/SAR detections (correlated and uncorrelated) and the uncorrelated interpolated reported positions.



**Figure 7.** Some types of false alarms in the example image. A blue rectangle is a detection in VV, a green triangle in VH. Except for (f), the VV channel is shown. (a) Coastal infrastructure outside the land mask (green line). (b) Aquaculture structures. (c) Azimuth ambiguity of a ship that is outside the image. The image edge can be seen at the bottom. (d) Non-AIS-reporting oil platform. (e,f) Detected ship plus false alarm on the cross-pol signature of the ship's wake ((e) VV image; (f) VH image). Orientation as in Figure 5, i.e., azimuth direction upward, range to the right.



**Figure 8.** Area in front of the port of Tarragona, VV channel. Almost all of the echoes on the water are due to azimuth ambiguities from land reflectors. The ones that are bright enough to have been detected by SUMO are indicated with a blue rectangle (detected in VV) and/or a green triangle (detected in VH); but the red dots indicate the ones that were automatically recognized as azimuth ambiguities. The only three detected real targets were detected in cross-pol and are a buoy (center) and two ships (top right). Orientation as in the previous figure.

## 6. Performance and Accuracy

One has to distinguish between the performance of the SAR imaging and that of a particular detector. The former poses the dominant limitations.

Performance of the SAR imaging in terms of actual detection rate (how many ships are present in the imaged area have been detected) and actual false alarm rate (how many objects have been detected that are not ships) is very difficult if not impossible to evaluate and validate. This is because in the first place, these figures depend on many factors, related to the targets (ship size, ship material, etc.), the sensor (resolution, polarization, etc.), the environment (wind, waves, etc.) and the imaging geometry (incidence angle, aspect angles, etc.). Detection and false alarm rates should be quantified as a function of all of these variables. In the second place, even in dedicated trials where test targets are deployed at sea, one can never be sure about the ground truth; in particular, one can never know if maybe some unknown small ships were present in the  $100 \text{ km} \times 100 \text{ km}$  image area. Nonetheless, some early quantitative performance figures have been published in [27,42] (using early versions of SUMO), and some recent results can be found in [80,81] (using several different detectors).

For the same reasons (dependency on many factors), from the point of view of the SAR imaging (regardless of the detector used), it is difficult to specify a minimum detectable ship size. In fact, this is an ill-defined concept. Large non-metallic ships can have a very small radar echo, while on the other hand, under low wind conditions, objects much smaller than the image resolution can be easily detected, such as small buoys fitted with radar reflectors. For common metallic ships, a very rough indication can be that the minimum detectable size is about half the resolution (i.e., equal to the pixel size). Some more detailed assessments can be found, e.g., in [82,83].

Regarding the performance of the detection algorithm, comparison with other ship detectors has been done in the past, finding a good level of performance for SUMO w.r.t. the others and also identifying many issues that the various detectors have in common [37]. A more recent benchmarking was done in [40]. Some extensive comparisons between ships detected in SAR images and ships that are self-reporting have been done recently with SUMO [35,36,39], but their outcomes reflect more the properties of the ship traffic, as well as being influenced by the images used besides the SUMO performance. No recent extensive quantitative evaluation of SUMO's performance has been done.

Further insight into the detector performance can be obtained by the visual validation of the automatic detection results. The authors know of no automatic detection algorithm that outperforms visual analysis by a human expert (except of course on speed). Regarding detection rate, the experience shows that a human expert can only improve marginally on the automatic SUMO algorithm results. A human can find ships that are not automatically detected because they fall inside the buffered land mask. In some cases, he may be able to discern the presence of two (or more) ships that are very close, but were detected as a single target. As discussed, the CFAR approach will not detect the smallest targets in order to limit the number of false alarms, and while this is a fundamental performance limitation of radar images, a human operator can use information and knowledge not available to the automatic detector (e.g., from context) to accept some targets below the formal detection threshold.

False alarms, on the other hand, are where a human operator can still much improve on the automatic SUMO results. The SAR signature (shape) of false alarms can look to an experienced operator tellingly different from that of a real ship, and no advanced recognition algorithms have been implemented in SUMO to deal with that. Past experience has shown that it is difficult to design such algorithms independently of sensor resolution. Furthermore, false alarms can be recognized by a human operator from the context: e.g., to expect reefs around tropical islands, or icebergs near a calving glacier in the Arctic, or to notice that a detection falls on a bright streak that is part of a pattern of streaks, e.g., made by an internal wave. The main false alarm causes for SUMO (and indeed for SAR ship detection algorithms in general) that can in many cases be recognized as such by an experienced human operator are:

- Unmasked islands and reefs,
- Strong waves,
- Ocean/atmosphere features such as fronts, rain cells, internal waves,
- Coastal infrastructure,
- Ship wakes,
- Edges of zero-wind areas,
- Icebergs,
- Range ambiguities, azimuth ambiguities with a source outside the image, radio interference.

The first three are more readily recognized if part of a wider pattern and less if isolated occurrences. The last two are also often very difficult to recognize for a human operator. Many of these were already noted in the DECLIMS project (Detection and Classification of Marine Traffic from Space [37]). For all of these false alarms, the computed reliability figure in SUMO does not correctly reflect the likelihood of a false alarm.

On account of this, the operator validation after the automatic detection is still valuable. On the other hand, in fully automatic mode without operator validation, the detection threshold parameter needs to be increased, especially in the co-pol channels, to prevent too many false alarms [39]. Therefore, the fully automatic mode is less sensitive and relies mostly on the ship echo in the cross-pol data.

Regarding the attributes of the detection, a quantitative evaluation of the accuracy of the ship size estimates by SUMO on Sentinel-1 images was published recently [72]. This evaluation used AIS data for reference and consequently did not include very small targets. The SAR imaging can introduce a blurring of the targets in the azimuth direction due to their motions (rotations and accelerations) on the waves. The blurring can be of the order of 100 m [84]. This can lead to serious overestimations of the ship size especially for small targets. Any extension of the signature in azimuth direction can be due to such blurring instead of reflecting the outline of the ship. A heading within  $5^{\circ}$ – $10^{\circ}$  of azimuth direction may indicate azimuth blurring instead of the real heading. Furthermore, in cross-pol images, the ship's wake is often seen as an extended bright region at the actual location of the ship (not displaced by the azimuth shift that acts on the moving ship) [75]. If this bright region is isolated from the displaced ship signature, it may give rise to a false alarm. If the two are connected, the ship's attributes may be wrong.

In the current Java software implementation of SUMO, as regards processing time, the data ingestion and the detection process at the pixel level are fast. The clustering is slower, so that the total processing time of an image depends more on the number of targets in the image than on its size. Furthermore, the land masking is a significant factor in the case of very convoluted coastlines. To process a dual-pol Sentinel-1 image of IW-GRDH-mode with 50 targets and a simple coastline takes about four minutes on an Intel Core2 Quad CPU at 2.83 GHz, 8 GB RAM, 64-bit Windows OS PC. In a recent big data trial, 11,400 such images with on average 55 targets per image were processed by SUMO in batch mode on a 480-core Linux cluster in 78 minutes, meaning a throughput of less than 0.5 s per image.

The main known shortfalls of SUMO that have not been mentioned so far are:

- ScanSAR MGD images are normally mosaicked from different sub-swaths. Each sub-swath may have a different ENL and PRF. SUMO however only uses a single value of the ENL over the whole image. It is adapted to the lowest sub-swath ENL and, therefore, not optimal for the others. SUMO can use a different PRF per sub-swath if the sub-swath boundaries are defined in the metadata. However, for several sensors, adjacent sub-swaths have an overlap zone that is computed as an average. Although this does lead to a lower speckle noise level in the overlap zone, which is good, it also introduces a doubling of the azimuth ambiguities, one set of ambiguities from each PRF in the overlap zone. This can give rise to errors.
- Although ship wakes are recognized to be a source of information on the ship traffic, SUMO does not analyze them.

- SUMO was designed for and on satellite SAR images. Tests with satellite optical images have shown that it can perform on those under favorable circumstances. However, common issues in optical images, such as clouds, whitecaps or sun glitter, are not dealt with. No tests with airborne radar images have been done, but presumably it should be possible to analyze such images at resolutions and incidence angles comparable to those of the satellite SARs.

## 7. Conclusions

The SUMO ship detector algorithm for satellite radar images is based on the classic and straightforward CFAR approach of detecting pixels that are high in relation to their local background. It combines a rigorous approach with a stochastic model for the background clutter PDF and an imposed false alarm rate with ad-hoc adjustments to deal with the deviations between model and reality and is designed with an eye on fast operation. It has proven to work across a wide range of image types: frequency bands L, C, X; Spotlight, Stripmap and ScanSAR modes; resolutions of 1 m–100 m; co- and cross-polarizations; SLC or MGD. It can detect a wide variety of maritime targets and estimate their properties, from 1 m-sized buoys to 450 m-sized super tankers, under a wide range of environmental conditions, always limited of course by the physical boundaries of radar observation.

Used fully automatically, its Java implementation is fast enough to keep up with current high production volumes of Sentinel-1, although then it depends for the most part on cross-pol data, as the co-pol images produce too many false alarms for fully automatic bulk processing. Used semi-automatically, it requires minimum time from a human operator to weed out residual false alarms, and then, the results are close to the best accuracy and completeness that can be expected from radar images, considering their inherent limited discriminative power.

**Acknowledgments:** Apart from the authors, many people have contributed to SUMO over time: Iain Shepherd, Guillermo Schwartz, Guido Lemoine, Tony Bauna, Jolyon Chesworth, Elena Aresu, Cristian Cesena, Marte Indregard, Thomas Lefort, Andrea Gabban, Juan Cicuendez. In Section 5, the image used is © Copernicus 2016; the AIS data used are from the Maritime Safety and Security Information System (MSSIS) terrestrial AIS network, courtesy of the Volpe Center of the U.S. Department of Transportation, the U.S. Navy and SPAWAR; and the coastline data are derived from OpenStreetMap, © OpenStreetMap contributors.

**Author Contributions:** Kourti initiated the developments and made the first SUMO algorithm. Alvarez, Santamaria and Greidanus developed the algorithm further and performed extensive testing and validation. Greidanus wrote this paper. Thoorens and Argentieri implemented the algorithm versions in software.

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

## Abbreviations

AIS	Automatic Identification System
CFAR	Constant False Alarm Rate
DEM	Digital Elevation Model
ENL	Equivalent Number of Looks
ESA	European Space Agency
EU	European Union
JRC	Joint Research Centre
LRIT	Long-Range Identification and Tracking
MGD	Multi-look Ground-range Detected
NOAA	U.S. National Oceanographic and Atmospheric Administration
PDF	Probability Density Function
PoI	Pixel of Interest
PRF	Pulse Repetition Frequency
RCS	Radar Cross-Section
SAR	Synthetic Aperture Radar
SLC	Single-Look Complex
VDS	Vessel Detection System
VMS	Vessel Management System



## References

1. Ulaby, F.T.; Moore, R.K.; Fung, A.K. *Microwave Remote Sensing: Active and Passive, Volume II: Radar Remote Sensing and Surface Scattering and Emission Theory*; Artech House: Norwood, MA, USA, 1986.
2. McCandless, S.W., Jr.; Jackson, C.R. Principles of Synthetic Aperture Radar. In *Synthetic Aperture Radar Marine User's Manual*; Jackson, C.R., Apel, J.R., Eds.; Department of Commerce, National Oceanic and Atmospheric Administration: Washington, DC, USA, 2004; pp. 1–24.
3. Ulaby, F.T.; Moore, R.K.; Fung, A.K. *Microwave Remote Sensing: Active and Passive, Volume III: From Theory to Applications*; Artech House: Norwood, MA, USA, 1986.
4. Holt, B. SAR imaging of the ocean surface. In *Synthetic Aperture Radar Marine User's Manual*; Jackson, C.R., Apel, J.R., Eds.; Department of Commerce, National Oceanic and Atmospheric Administration: Washington, DC, USA, 2004; pp. 25–80.
5. Alpers, W. C-band radar signatures of rain over the ocean-revisited. In Proceedings of the IGARSS 2015, Milan, Italy, 26–31 July 2015.
6. Danklmayer, A.; Döring, B.J.; Schwerdt, M.; Madhu, C. Assessment of atmospheric propagation effects in SAR Images. *IEEE Trans. Geosci. Rem. Sens.* **2009**, *47*, 3507–3518. [[CrossRef](#)]
7. Pichel, W.G.; Clemente-Colon, P.; Wackerman, C.C.; Friedman, K.S. Ship and wake detection. In *Synthetic Aperture Radar Marine User's Manual*; Jackson, C.R., Apel, J.R., Eds.; Department of Commerce, National Oceanic and Atmospheric Administration: Washington, DC, USA, 2004; pp. 277–303.
8. Crisp, D.J. *The State-of-the-Art in Ship Detection in Synthetic Aperture Radar Imagery*; Defence Science and Technology Organisation, Information Sciences Laboratory: Edinburgh, South Australia, 2004.
9. Arnesen, T.N.; Olsen, R.B. *Literature Review on Vessel Detection*; Norwegian Defence Research Establishment (FFI): Kjeller, Norway, 2004.
10. Vachon, P.W.; Campbell, J.W.M.; Bjerkelund, C.A.; Dobson, F.W.; Rey, M.T. Ship detection by the RADARSAT SAR: Validation of detection model predictions. *Can. J. Remote Sens.* **1997**, *23*, 48–59. [[CrossRef](#)]
11. Wackerman, C.C.; Friedman, K.S.; Pichel, W.G.; Clemente-Colón, P.; Li, X. Automatic detection of ships in RADARSAT-1 SAR imagery. *Can. J. Remote Sens.* **2001**, *27*, 371–378. [[CrossRef](#)]
12. Eldhuset, K. An automatic ship and ship wake detection system for spaceborne SAR images in coastal regions. *IEEE Trans. Geosci. Rem. Sens.* **1996**, *34*, 1010–1019. [[CrossRef](#)]
13. Tello, M.; López-Martínez, C.; Mallorqui, J.J. A novel algorithm for ship detection in SAR imagery based on the wavelet transform. *IEEE Geosci. Rem. Sens.* **2005**, *2*, 201–205. [[CrossRef](#)]
14. Leung, H.; Dubash, N.; Xie, N. Detection of small objects in clutter using a GA-RBF neural network. *IEEE Trans. Aerosp. Electr. Syst.* **2002**, *38*, 98–118. [[CrossRef](#)]
15. Martín-de-Nicolás, J.; Mata-Moya, D.; Jarabo-Amores, M.P.; del-Rey-Maestre, N.; Bárcena-Humanes, J.L. Neural network based solutions for ship detection in SAR images. In Proceedings of the 18th International Conference on Digital Signal Processing, Fira, Greece, 1–3 July 2013.
16. Lombardo, P.; Sciotti, M. Segmentation-based technique for ship detection in SAR images. *IEEE Proc. Radar Sonar Navig.* **2001**, *148*, 147–159. [[CrossRef](#)]
17. Gao, G. A Parzen-Window-Kernel-Based CFAR algorithm for ship detection in SAR Images. *IEEE Geosci. Remote Sens. Lett.* **2011**, *8*, 557–561. [[CrossRef](#)]
18. Yu, Y.; Bi, F.; Pu, Y.; Zhou, H.; Yang, J. Saliency-based ship detection in SAR images. In Proceedings of the IET International Radar Conference 2015, Hangzhou, China, 14–16 October 2015.
19. Ouchi, K.; Tamaki, S.; Yaguchi, H.; Iehara, M. Ship detection based on coherence images derived from cross correlation of multilook SAR images. *IEEE Geosci. Remote Sens. Lett.* **2004**, *1*, 184–187. [[CrossRef](#)]
20. Gambardella, A.; Nunziata, F.; Migliaccio, M. A Physical full-resolution SAR ship detection filter. *IEEE Geosci. Rem. Sens. Lett.* **2008**, *5*, 760–763. [[CrossRef](#)]
21. Y Jeremy, M.; Campbell, J.W.M.; Mattar, K.; Potter, T. Ocean surveillance with polarimetric SAR. *Can. J. Remote Sens.* **2001**, *27*, 328–344. [[CrossRef](#)]
22. Touzi, R.; Charbonneau, F.J.; Hawkins, R.K.; Vachon, P.W. Ship detection and characterization using polarimetric SAR. *Can. J. Rem. Sens.* **2004**, *30*, 552–559. [[CrossRef](#)]
23. Nunziata, F.; Migliaccio, M.; Brown, C.E. Reflection symmetry for polarimetric observation of man-made metallic targets at sea. *IEEE J. Ocean. Eng.* **2012**, *37*, 384–394. [[CrossRef](#)]

24. Marino, A.; Sugimoto, M.; Ouchi, K.; Hajnsek, I. Validating a notch filter for detection of targets at sea with ALOS-PALSAR Data: Tokyo Bay. *IEEE J. Sel. Top. Appl. Earth Obs. Rem. Sens.* **2014**, *7*, 4907–4918. [[CrossRef](#)]
25. Touzi, R.; Hurley, J.; Vachon, P.W. Optimization of the degree of polarization for enhanced ship detection using polarimetric RADARSAT-2. *IEEE Trans. Geosci. Remote Sens.* **2015**, *53*, 5403–5424. [[CrossRef](#)]
26. Lemoine, G.; Schwartz-Juste, G.; Kourti, N.; Shepherd, I.; Cesena, C. An open source framework for integration of vessel positions detected in spaceborne SAR Imagery in operational fisheries monitoring and control. In Proceedings of the 2004 Envisat & ERS Symposium, Salzburg, Austria, 6–10 September 2004.
27. Kourti, N.; Shepherd, I.; Greidanus, H.; Alvarez, M.; Aresu, E.; Bauna, T.; Chesworth, J.; Lemoine, G.; Schwartz, G. Integrating remote sensing in fisheries control. *Fish. Manag. Ecol.* **2005**, *12*, 295–307. [[CrossRef](#)]
28. Lemoine, G.G.; Greidanus, H.; Shepherd, I.M.; Kourti, N. Developments in satellite fisheries monitoring and control. In Proceedings of the 8th International Conference on Remote Sensing for Marine and Coastal Environments, Halifax, NS, Canada, 17–19 May 2005.
29. Lefort, T.; Thoorens, F.; Greidanus, H.; Indregard, M.; Gallagher, R.; Aznar-Abian, V.; Defever, S.; Ramponi, N. Providing cost efficient near real time fisheries monitoring capability to remote and isolated areas of the globe—The case of Comoros. In Proceedings of the 33rd International Symposium on Remote Sensing of Environment, Stresa, Italy, 4–8 May 2009; pp. 833–836.
30. Cicuendez Perez, J.; Alvarez Alvarez, M.; Heikkonen, J.; Indregard, M. Identification of bluefin tuna cages in Mediterranean Sea fishing grounds from SAR images. *Int. J. Rem. Sens.* **2011**, *32*, 4461–4474. [[CrossRef](#)]
31. Cicuendez Perez, J.; Alvarez Alvarez, M.; Heikkonen, J.; Guillen, J.; Barbas, T. The efficiency of using remote sensing for fisheries enforcement: Application to the Mediterranean bluefin tuna fishery. *Fish. Res.* **2003**, *147*, 24–31.
32. Greidanus, H.; Ardill, D.; Ansell, N.; Thoorens, F.X.; Peries, P. Surveying shipping and fishing in the SW Indian Ocean with satellite SAR. In Proceedings of the 33rd International Symposium on Remote Sensing of Environment, Stresa, Italy, 4–8 May 2009.
33. Posada, M.; Greidanus, H.; Alvarez, M.; Vespe, M.; Cokacar, T.; Falchetti, S. Maritime awareness for counter-piracy in the Gulf of Aden. In Proceedings of IGARSS 2011, Vancouver, ON, Canada, 24–29 July 2011; pp. 249–252.
34. Greidanus, H.; Alvarez, M.; Eriksen, T.; Argentieri, P.; Cokacar, T.; Pesaresi, A.; Falchetti, S.; Nappo, D.; Mazzarella, F.; Alessandrini, A. Basin-wide maritime awareness from multi-source ship reporting data. *TransNav Int. J. Mar. Navig. Saf. Sea Transp.* **2013**, *7*, 35–42. [[CrossRef](#)]
35. Greidanus, H.; Santamaria, C.; Alvarez, M.; Krause, D.; Stasolla, M.; Vachon, P.W. Non-Reporting ship traffic in the western Indian Ocean. In Proceedings of the ESA Living Planet Symposium, Prague, Czech Republic, 9–13 May 2016.
36. Santamaria, C.; Greidanus, H.; Fournier, M.; Eriksen, T.; Vespe, M.; Alvarez, M.; Fernandez Arguedas, V.; Delaney, C.; Argentieri, P. Sentinel-1 contribution to monitoring maritime activity in the arctic. In Proceedings of the ESA Living Planet Symposium, Prague, Czech Republic, 9–13 May 2016.
37. Greidanus, H.; Clayton, P.; Indregard, M.; Staples, G.; Suzuki, N.; Vachon, P.; Wackerman, C.; Tennvassas, T.; Mallorqui, J.; Kourti, N.; et al. Benchmarking operational SAR ship detection. In Proceedings of the IGARSS 2004, Anchorage, AL, USA, 20–24 September 2004.
38. Greidanus, H.; Kourti, N. Findings of the DECLIMS project—Detection and classification of marine traffic from space. In Proceedings of the SEASAR 2006: Advances in SAR oceanography from ENVISAT and ERS, ESA-ESRIN, Frascati, Italy, 23–26 January 2006.
39. Santamaria, C.; Stasolla, M.; Fernandez Arguedas, V.; Argentieri, P.; Alvarez, M.; Greidanus, H. *Sentinel-1 Maritime Surveillance-Testing and Experiences with Long-Term Monitoring*; JRC Science and Policy Report; Publications Office of the European Union: Luxembourg, 2015.
40. Stasolla, M.; Mallorqui, J.J.; Margarit, G.; Santamaria, C.; Walker, N. A comparative study of operational vessel detectors for maritime surveillance using satellite-borne Synthetic Aperture Radar. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2016**, *9*, 2687–2701. [[CrossRef](#)]
41. Council Regulation (EC). No 1224/2009 of 20 Nov 2009 Establishing a Community Control System for Ensuring Compliance with the Rules of the Common Fisheries Policy; European Commission: Brussels, Belgium, 2009.
42. Kourti, N.; Shepherd, I.; Schwartz, G.; Pavlakakis, P. Integrating spaceborne SAR imagery into operational systems for fisheries monitoring. *Can. J. Remote Sens.* **2001**, *27*, 291–305. [[CrossRef](#)]

43. Ward, K.D. Compound representation of high resolution sea clutter. *Electron. Lett.* **1981**, *17*, 561–563. [[CrossRef](#)]
44. Oliver, C.; Quegan, S. *Understanding Synthetic Aperture Radar Images*; Artech House: Norwood, MA, USA, 1998.
45. Greidanus, H. Applicability of the K distribution to RADARSAT maritime imagery. In Proceedings of the IGARSS 2004, Anchorage, AL, USA, 20–24 September 2004; pp. 4715–4718.
46. Vespe, M.; Greidanus, H. SAR image quality assessment and indicators for vessel and oil spill detection. *IEEE Trans. Geosc. Remote Sens.* **2012**, *50*, 4726–4734. [[CrossRef](#)]
47. Gao, G. Statistical modeling of SAR images: A survey. *Sensors* **2010**, *10*, 775–795. [[CrossRef](#)] [[PubMed](#)]
48. Martín-de-Nicolás, J.; Jarabo-Amores, M.-P.; Mata-Moya, D.; del-Rey-Maestre, N.; Bárcena-Humanes, J.-L. Statistical analysis of SAR sea clutter for classification purposes. *Remote Sens.* **2014**, *6*, 9379–9411. [[CrossRef](#)]
49. Trunk, G.V.; George, S. F. Detection of targets in non-Gaussian sea clutter. *IEEE Trans. Aerosp. Electron. Syst.* **1970**, *6*, 620–628. [[CrossRef](#)]
50. Goldstein, G. False-alarm regulation in log-normal and Weibull clutter. *IEEE Trans. Aerosp. Electron. Syst.* **1973**, *9*, 84–92. [[CrossRef](#)]
51. Lorenzetti, J.A.; Paes, R.L.; Gherardi, D.M. A performance comparison of a CFAR ship detection algorithm using Envisat, Radarsat, Cosmo-Skymed and Terrasar-X images. In Proceedings SeaSAR 2010, Frascati, Italy, 25–29 January 2010.
52. Jakeman, E.; Pusey, P. A model for non-Rayleigh sea echo. *IEEE Trans. Antennas Propag.* **1976**, *24*, 806–814. [[CrossRef](#)]
53. Qin, X.; Zhou, S.; Zou, H.; Gao, G. Statistical modeling of sea clutter in high-resolution SAR images using generalized gamma distribution. In Proceedings 2012 International Conference on Computer Vision in Remote Sensing, Xiamen, China, 16–18 December 2012.
54. Martín-de-Nicolás, J.; Jarabo-Amores, P.; del-Rey-Maestre, N.; Gómez-del-Hoyo, P.; Bárcena-Humanes, J.L. Robustness of a generalized gamma CFAR ship detector applied to TerraSAR-X and Sentinel-1 images. In Proceedings IEEE International Conference Computer Tool (EUROCON), Salamanca, Spain, 8–11 September 2015.
55. Watts, S. Radar detection prediction in k-distributed sea clutter and thermal noise. *IEEE Trans. Aerosp. Electr. Syst.* **1987**, *23*, 40–45. [[CrossRef](#)]
56. Anastassopoulos, V.; Lampropoulos, G.A.; Drosopoulos, A.; Rey, M. High resolution radar clutter statistics. *IEEE Trans. Aerosp. Electr. Syst.* **1999**, *35*, 43–60. [[CrossRef](#)]
57. Li, F.K.; Johnson, W.T.K. Ambiguities in spaceborne Synthetic Aperture Radar systems. *IEEE Trans. Aerosp. Electr. Syst.* **1983**, *19*, 389–397. [[CrossRef](#)]
58. Olsen, R.B.; Arnesen, T.N.; Eldhuset, K. Signatures of vessels in ENVISAT AP-mode imagery. In Proceedings of the IGARSS 2004, Anchorage, AL, USA, 20–24 September 2004; pp. 3895–3897.
59. Souyris, J.; Imbo, P.; Fjørtoft, R.; Mingot, S.; Lee, J. compact polarimetry based on symmetry properties of geophysical media: The  $\pi/4$  Mode. *IEEE Trans. Geosci. Remote Sens.* **2005**, *43*, 643–646. [[CrossRef](#)]
60. Raney, R.K. Hybrid-Polarity SAR architecture. *IEEE Trans. Geosc. Remote Sens.* **2007**, *45*, 3397–3404. [[CrossRef](#)]
61. Shirvany, R.; Chabert, M.; Tournet, J. Ship and oil-spill detection using the degree of polarization in linear and hybrid/compact Dual-Pol SAR. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2012**, *5*, 885–892. [[CrossRef](#)]
62. Smith, A.J.E.; Chesworth, J.; Greidanus, H. Ship detection with Envisat’s alternating polarisation mode. In Proceedings of the 2004 Envisat & ERS Symposium, Salzburg, Austria, 6–10 September 2004.
63. Arnesen, H.; Tonje, N. Multi-channel and multi-polarisation ship detection. In Proceedings of the IGARSS 2012, Munich, Germany, 22–27 July 2012.
64. Gao, G.; Shi, G.; Zhou, S. Ship detection in high-resolution dual-polarization SAR amplitude images. *Int. J. Antennas Propag.* **2013**. [[CrossRef](#)]
65. Theoharatos, C.; Makedonas, A.; Fragoulis, N.; Tsagaris, V.; Costicoglou, S. Detection of ship targets in polarimetric SAR data using 2D-PCA data fusion. In Proceedings of the 36th International Symposium on Remote Sensing of Environment, Berlin, Germany, 11–15 May 2015.
66. Liu, C. *A Dual-Polarization Ship Detection Algorithm*; Defence Research and Development Canada, Ottawa Research Centre: Ottawa, ON, Canada, 2015.

67. Wessel, P.; Smith, W.H.F. A global self-consistent, hierarchical, high-resolution shoreline database. *J. Geophys. Res.* **1996**, *101*, 8741–8743. [CrossRef]
68. OpenStreetMap. Available online: <https://www.openstreetmap.org/> (accessed on 18 November 2016).
69. U.S. National Oceanic and Atmospheric Administration–National Ice Center/U.S. Navy–Naval Ice Centre. Available online: [http://www.natice.noaa.gov/products/daily\\_products.html](http://www.natice.noaa.gov/products/daily_products.html) (accessed on 18 November 2016).
70. Blacknell, D. Comparison of parameter estimators for K-distribution. *IEEE Proc. Radar Sonar Navig.* **1994**, *141*, 45–52. [CrossRef]
71. Kaiser, J.F. Nonrecursive digital filter design using the 10-sinh window function. In Proceedings of the 1974 IEEE International Symposium on Circuits and Systems, San Francisco, CA, USA, 20–23 April 1974.
72. Stasolla, M.; Greidanus, H. The exploitation of Sentinel-1 images for vessel size estimation. *Remote Sens. Lett.* **2016**, *7*, 1219–1228. [CrossRef]
73. Knott, E.F. Radar cross section. In *Radar Handbook*, 2nd ed.; McGraw-Hill: Boston, MA, USA, 1990; pp. 11.1–11.51.
74. Greidanus, H.; Santamaria, C. *First Analyses of Sentinel-1 Images for Maritime Surveillance*; Publications Office of the European Union: Luxembourg, 2014.
75. Greidanus, H. Satellite imaging for maritime surveillance of the European seas. In *Remote Sensing of the European Seas*; Barale, V., Gade, M., Eds.; Springer: Berlin, Germany, 2008; pp. 343–358.
76. Meadows, P. *S1-A Annual Performance Report for 2015*; ESA (European Space Agency): San Francisco, CA, USA, 2016.
77. Mazzarella, F.; Alessandrini, A.; Greidanus, H.; Alvarez, M.; Argentieri, P.; Nappo, D.; Lukasz, Z. Data fusion for wide-area maritime surveillance. In Proceedings of the COST MOVE Workshop on Moving Objects at Sea, Brest, France, 27–28 June 2013.
78. Santamaria, C.; Greidanus, H. Ambiguity discrimination for ship detection using Sentinel-1 repeat acquisition operations. In Proceedings of the IGARSS 2015, Milan, Italy, 26–31 July 2015.
79. European Space Agency. *Sentinel-1 Product Definition. Document S1-RS-MDA-52-7440*; ESA (European Space Agency): San Francisco, CA, USA, 2016.
80. Pelich, R.; Longépé, N.; Mercier, G.; Hajduch, G.; Garello, R. Performance evaluation of Sentinel-1 data in SAR ship detection. In Proceedings of the 2015 IEEE International Geoscience and Remote Sensing Symposium, Milan, Italy, 26–31 July 2015.
81. Pelich, R.; Longépé, N.; Mercier, G.; Hajduch, G.; Garello, R. AIS-Based evaluation of target detectors and SAR sensors characteristics for maritime surveillance. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2015**, *8*, 3892–3901. [CrossRef]
82. Vachon, P.W.; Wolfe, J. *GMES Sentinel-1 Analysis of Marine Applications Potential (AMAP)*; DRDC Ottawa ECR2008-218, Defence Research and Development Canada: Ottawa, ON, Canada, 2008.
83. Vachon, P.W.; Wolfe, J.; Greidanus, H. Analysis of Sentinel-1 marine applications potential. In Proceedings of the 2012 IEEE International Geoscience and Remote Sensing Symposium, Munich, Germany, 22–27 July 2012.
84. Greidanus, H. Sub-aperture behaviour of SAR signatures of ships. In Proceedings of the IGARSS 2006, Denver, CO, USA, 31 July–4 August 2006.



© 2017 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).