European Commission

JRC SCIENCE AND POLICY REPORTS

Sentinel-1 Maritime Surveillance

Testing and Experiences with Long-term Monitoring

Carlos Santamaria, Mattia Stasolla,
Virginia Fernandez Arguedas,
Pietro Argentieri, Marlene Alvarez,
Harm Greidanus

2015

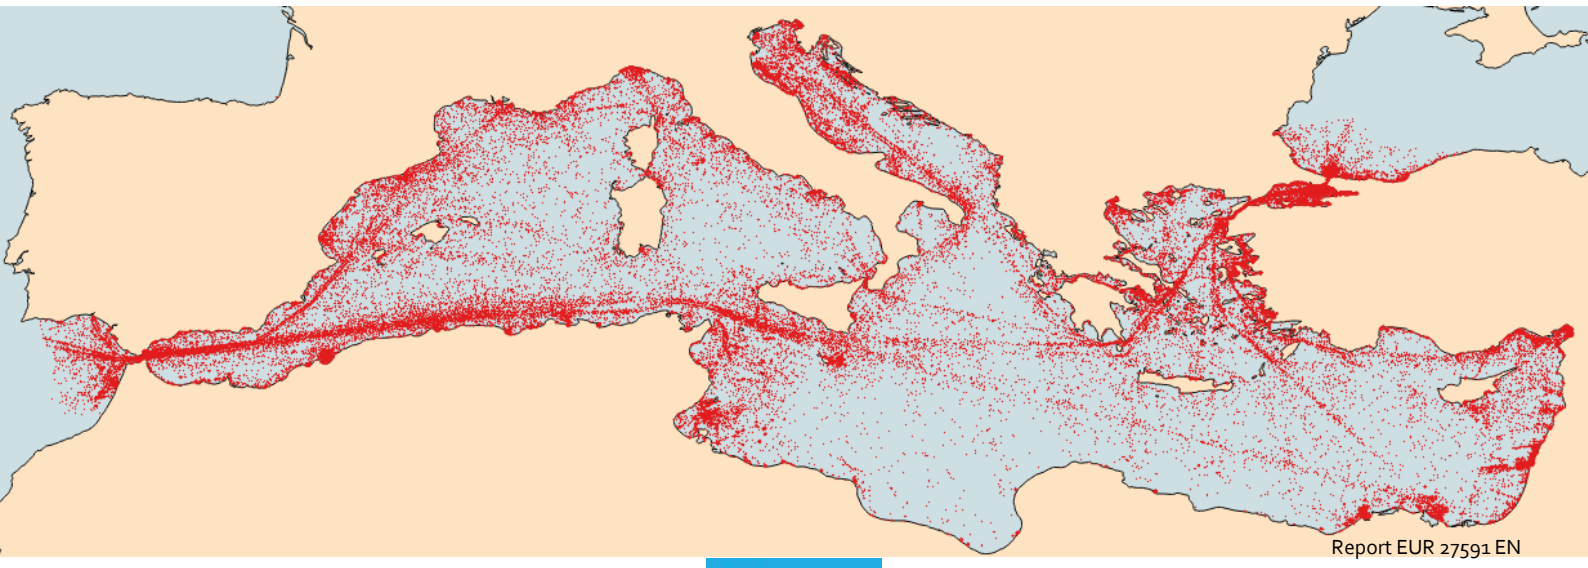

Report EUR 27591 EN

Joint Research Centre

European Commission
Joint Research Centre
Institute for the Protection and Security of the Citizen

Contact information
Harm Greidanus
Address: Joint Research Centre, TP670, Via E. Fermi 2749, 21027 Ispra, Italy
E-mail: harm.greidanus@jrc.ec.europa.eu
Tel.: +39 0332 78 9739

https://ec.europa.eu/jrc/

JRC98532

EUR 27591 EN

ISBN 978-92-79-53960-2 (PDF)

ISSN 1831-9424 (online)

doi:10.2788/090400

Luxembourg: Publications Office of the European Union, 2015

Abstract

Sentinel-1 is the Synthetic Aperture Radar (SAR) satellite of the EU's Copernicus program for Earth observation. Being a radar, it is very suitable for ship detection. Sentinel-1 produces routinely a large amount of images over the seas and coasts, in particular in Europe and the Arctic but also elsewhere on the globe. Up to now, software to analyse satellite SAR images for ship detection was limited to handle a few images at the time, so was not up to the high production volume of Sentinel-1. This report describes improvements in JRC's ship detector SUMO to fully automatically process large amounts of images. With this increased capability, time series of many months of Sentinel-1 images are analysed over two sites in the Arctic, revealing the seasonal relation between ship traffic and sea ice. Also a multi-month data set over the Mediterranean Sea, mapping the shipping patterns there, and a one-year data set over the Western Indian Ocean are analysed. The combination of the Sentinel-1 ship detections with ship positions derived from ship self-reporting systems (AIS and LRIT) results in the quantification and mapping of the non-reporting ship traffic, that is not seen on AIS or LRIT. Beyond mere detection, new methods are tested to estimate the size of the ships from their signature in the Sentinel-1 images, and to allow discrimination between tankers and cargo vessels even if they are of similar size. The occurrence of range ambiguities, and the impact of SAR imaging effects on the apparent heading of the ship signatures are briefly discussed.
The developments and experiences discussed here should further promote Sentinel-1 in becoming a tool for maritime surveillance to the benefit of maritime safety, security and sustainability, including spatial planning.

Front cover: 92,019 maritime target detections from 1,557 Sentinel-1 images over the Mediterranean Sea

Executive summary

The Copernicus program for space-based Earth observation represents a major investment of the EU. Its exploitation should be maximised. One application area of Copernicus is maritime surveillance: monitoring shipping activity at sea. Maritime surveillance is a tool for security, safety and sustainability. As an example in each of these three domains, by monitoring and understanding ship behaviour it becomes possible to be alerted to potential illegal border crossings, infringements of maritime traffic separation schemes, and illegal fishing. Maritime surveillance is also useful for maritime spatial planning including verifying the actual use of the maritime space, which are issues of increasing interest. Space-based sensors are eminently suitable for maritime surveillance because they reach beyond the coastal ranges, providing global, regularly repeated coverage. Of the Copernicus satellites, Sentinel-1 is the most suitable for maritime surveillance, because its Synthetic Aperture Radar (SAR) can detect ships, and it can do so independently of cloud cover or daylight. Also Sentinel-2 can detect ships with its optical imager, but not under a cloud cover or at night.

This report discusses the results and experiences from trials for ship detection and ship traffic monitoring with Sentinel-1, based on one year of data availability. The following topics are specifically addressed:

- Fully automatic ship detection in Sentinel-1 images;
- Long-term, wide-area monitoring of ship traffic;
- Applications for ship traffic monitoring in the Arctic;
- Ship size estimation and classification.

Automatic ship detection is an issue, because Sentinel-1 produces every day a large amount of images over the sea surface – too much to analyse by human operators. If these data are going to be used for maritime surveillance, fully automatic analysis methods are needed. Up to now, however, fully automatic ship detection did not produce reliable results; validation of the fully automatic results by a human operator is normally needed to weed out false alarms (radar detections that are not due to ships).

This report presents improvements on the automatic ship detection, as implemented in JRC's ship detector SUMO, that enable batch processing of large amounts of Sentinel-1 images without individual human operator assistance. The improvements include in particular: an improved geolocation of the images based on satellite orbit data that helps avoid errors in masking out the land; use of improved land masks; automatic rejection of Sentinel-1 image artefacts (azimuth ambiguities and edge effects); and flagging of recurrent targets in repeat-pass images.

The batch processing was applied to all Sentinel-1 images that have been collected in the Western Indian Ocean off East Africa during one year – a total of 1,874 images, yielding 21,520 ship detections. This exercise was done in relation to the PMAR (Piracy, Maritime Awareness and Risks) project that JRC carried out to support authorities in Africa with counter-piracy and maritime security, as that project provided access to a very complete set of ship self-reporting data (AIS and LRIT) that could be used for verification. Secondly, the batch processing was applied to several months worth of Sentinel-1 images over the Mediterranean Sea – 1,557 images yielding 92,019 ship detections. In combination with ship traffic data from ship reporting systems, this exercise has quantified the occurrence of non-reporting ships in both areas. In the Med Sea, where the coverage of the Sentinel-1 imagery was quite uniform, it has revealed the shipping patterns, to a higher level of completeness than what is obtained from the ship reporting systems as their coverage in the Med Sea is incomplete.

Having a polar orbit, Sentinel-1 passes over the Arctic very frequently, and is as such very suitable for Arctic surveillance. The results of multi-month monitoring of two test sites in the Arctic (northern Baltic and Kara Sea coast), using Sentinel-1 images in combination with satellite AIS, show how shipping reacts to the dynamic sea ice, including how once carved out ship tracks through the sea ice

can remain in place for much of the winter season. Challenges remain, however, in avoiding false alarms with ship detection in the presence of sea ice.

While satellite SAR is strong in the *detection* of ships (albeit with certain shortcomings), in *classification* (determining the ship type) it faces more difficulties. The first step in classification is size estimation. Classification and size estimation are difficult in SAR images because of limited resolution and because of SAR image distortions and artefacts that blur the outline and signature of moving targets. For Sentinel-1, the most frequently used products are IW-GRDH and EW-GRDM, with resolutions of 20 m and 90 m respectively. The latter has very limited ability for ship classification, a challenging task already for the 20 m resolution. The report includes a benchmarking exercise with four different algorithms to estimate the ship size from Sentinel-1 IW-GRDH images, using 107 known and manually verified ships. It is found that the most advanced algorithm, based on mathematical morphology, produces good length estimates and reasonable width estimates.

Classification ideally estimates the ship type from the signature in the SAR image. However, there are many ship types (tanker, container, bulk cargo, fishing, passenger, tug, military, etc.), and at 20 m resolution, the SAR signature is often not more than an elongated blob. A full classification is still beyond reach. Therefore, a reduced problem is considered here, namely the disambiguation between the two most occurring ship types that have a similar length, tankers and cargoes. (Disambiguation between e.g. a tug and a tanker can often be done based on length alone.) The same set of validated ships as used for the size estimation was used to test an algorithm for cargo-tanker disambiguation that is based on the texture in the SAR signature. By combining two texture measures, and splitting the ship signature in stern, middle and bow, a classification accuracy of 82 % could be obtained.

While the improvements implemented in the JRC's SUMO ship detection software and Blue Hub maritime surveillance development platform have enabled the automatic bulk processing of Sentinel-1 images for ship detection, there are still points that can be improved. For SUMO, the fully automatic mode today still has to use a higher detection threshold (favouring bigger ships) than the semi-automatic, operator-verified mode. Improved algorithms for the automatic discrimination of false alarms would enable to lower the detection threshold and consequently detect smaller ships. Land masks could be further improved to take into account features like reefs and tidal lands. Regarding ship size estimation and classification, the algorithms for that are at present run outside of SUMO; integration into SUMO would make this step more efficient. For automatic classification of ships in SAR images, much development can still be done. The analysis of recurrent targets, to suppress a wide category of false alarms that are due to radar echoes that are always found in the same place, is still in an early stage of development and can be refined. Regarding Arctic applications, the suppression / discrimination of sea ice in the ship detection should be developed and implemented if we want to be able to routinely perform maritime surveillance in Arctic winter waters. Finally, in JRC's Blue Hub, the algorithms for the automatic correlation between SAR ship detections and ship positions known from the ship reporting systems can be further improved, to provide more reliable indications of non-reporting ships. Development and implementation of these improvements will enable a wider exploitation of Sentinel-1 images.

Contents

1 Introduction 6
2 The SUMO automatic ship detector 9
2.1 SUMO 9
2.2 After SUMO – The Blue Hub 9
2.3 SUMO developments 10
3 Sentinel-1 in the Arctic 15
3.1 Yenisei Gulf 15
3.2 Bay of Bothnia 22
4 Sentinel-1 in the Western Indian Ocean 34
4.1 Ship detection in batch mode 36
5 Sentinel-1 in the Mediterranean Sea 42
6 Ship size estimation in Sentinel-1 images 46
6.1 Dataset 46
6.2 Results 48
6.3 Conclusions 53
7 Ship disambiguation in Sentinel-1 images 54
7.1 Dataset & ground truth 54
7.2 Cargo/tanker disambiguation based on visual analytics 55
7.3 Experimental results 58
7.4 Multi-feature cargo/tanker disambiguation 59
7.5 Results analysis and conclusions 60
8 Other issues 62
8.1 Duplicated ships in Sentinel-1 TOPS SLC images 62
8.2 Analysis of range ambiguities affecting maritime images 65
9 Conclusions 68
References 70
Acronyms 72
Acknowledgments 73

1 Introduction

Copernicus is the EU's flagship program on Earth observation, primarily intended for use by policymakers and public authorities in Europe but also for services to the wider economy and public [C2015a, C2015b]. Its main tools are a series of satellites, one of which is Sentinel-1, a spaceborne Synthetic Aperture Radar (SAR), operated by ESA [E2013, E2015]. Sentinel-1 produces images of the Earth surface, including the oceans, in radar wavelengths. As a radar, Sentinel-1 is suitable for ship detection, independent of day/night conditions and unhindered by clouds. It is therefore a useful tool for maritime surveillance, to monitor the ship traffic and to become aware of what is happening at sea.

Nowadays there are many tools available for maritime surveillance. They include on the one hand cooperative systems, in which the ships themselves report their identities and positions to others and to the authorities. The three most common automatic ship self-reporting systems are AIS, LRIT and VMS – the Automatic Identification System designed for collision avoidance, the Long Range Identification and Tracking designed for security, and the Vessel Monitoring System designed for fisheries control. Thanks to national, regional and global networks that share AIS data, and thanks to AIS-receiving satellites, AIS data can be available practically continuously and globally. LRIT and VMS data cover significant portions of the ship traffic, but their availability is severely restricted. On the other hand, there are the non-cooperative surveillance tools: observation systems that detect, track and/or identify ships without any help from the ship's side. These include e.g. cameras and radars, that can be located on the coast, on a ship, airplane or satellite. Sentinel-1 falls into this category.

In order to be well aware of what is happening at sea and to be in a position to recognise abnormal, risky or threatening situations, the data from the various systems need to be integrated and analysed. The JRC is performing research into this integration with its Blue Hub [BH2015], which is an R&D platform to exploit big data in the maritime domain and improve the understanding of what is happening at sea. Inasmuch as Sentinel-1 can be an important source of data for maritime awareness, for maximum value it should be exploited in fusion with other data sources, complementing in particular the ship self-reporting systems.

Satellite SARs like Sentinel-1 have been available since the 1990s. They circle the Earth with a revolution time of the order of 100 minutes, in nearly North-South orbits that always pass over the poles. As they pass over, they make (still) images of the Earth surface, with a level of detail (resolution) that can be as good as 1 m nowadays. However, such very detailed images can only be very small, about 5 km in size. Therefore, maritime surveillance generally prefers to use wider swath images that however have a lower resolution. Satellite SARs in existence today that can be used for maritime surveillance include TerraSAR-X (Germany), Cosmo-SkyMed (Italy), Radarsat-2 (Canada), Alos-Palsar-2 (Japan), Kompsat-5 (Korea) and Risat (India); Paz (Spain) is expected soon. ESA's predecessors of Sentinel-1 were ERS-1, ERS-2 and ENVISAT-ASAR.

What sets Sentinel-1 apart from all these other systems is that it is routinely collecting a large amount of images, as opposed to being tasked on demand; and the images are made freely available, as opposed to being commercially sold. See Figure 1-1 for an example of the area on Earth that Sentinel-1 routinely observes. This data acquisition policy leads to a very large amount of data being produced each day, many hundreds of images representing of the order of a Terabyte of data daily. Although most of the images are over land, a part are over seas and especially over coasts. Figure 1-1 is representative for the fact that the European seas, the Arctic seas and the North Atlantic are routinely covered. If all these data are going to be put to use for maritime surveillance, fully automatic processing is needed, which up to now was not available at a sufficient level of reliability / performance / robustness.

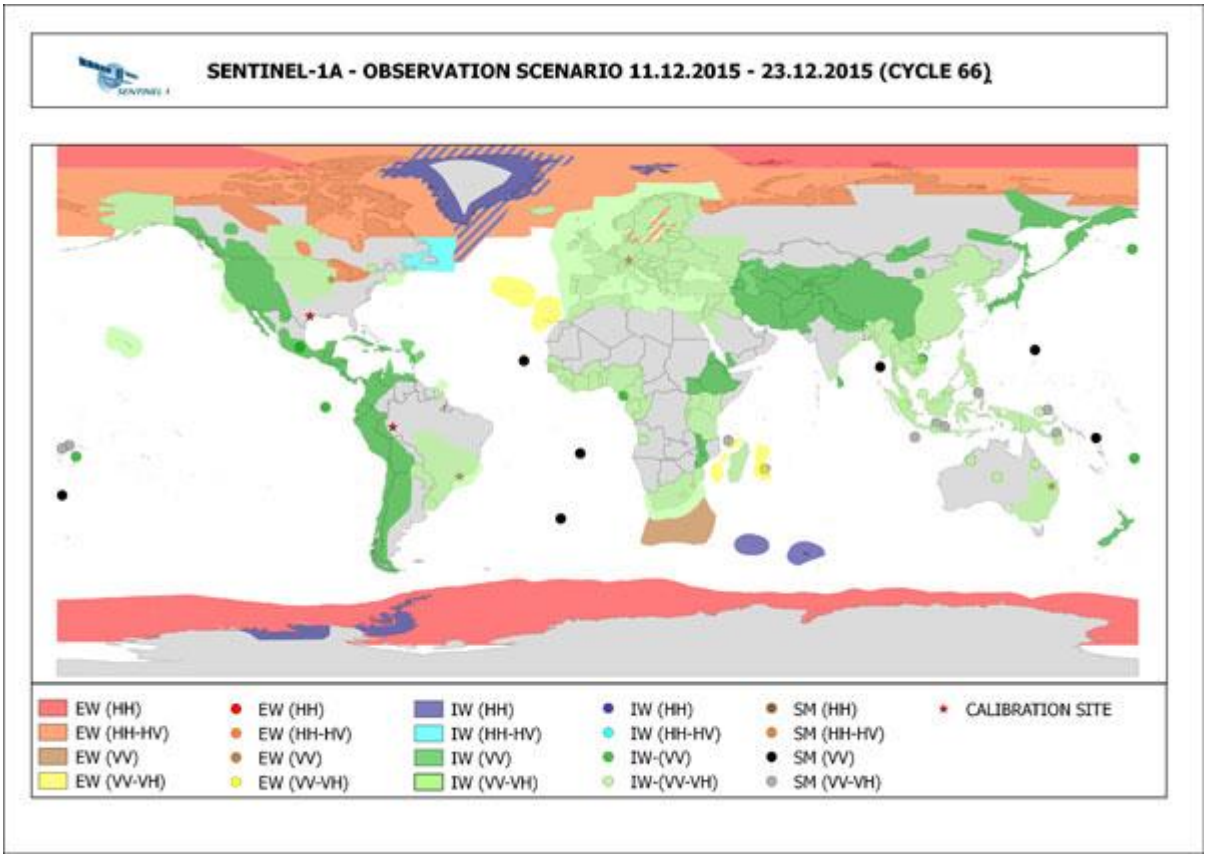

Figure 1-1. Sentinel-1's planned observation scenario for two weeks in December 2015. The different colours indicate different observation modes. (From ESA's website [E2015].)

Sentinel-1 was launched on 3 April 2014, and since 3 October 2014, its data are operationally available [E2014]. A prior report [GS2014] discussed the first experiences with its data. The current report documents the experiences with and results from using Sentinel-1 data at JRC during its first year of operations, between October 2014 and November 2015. First, the report discusses the progress made in the automation of the analysis of the Sentinel-1 images for ship detection (Chapter 2). Up to now, the ship detection processing was designed to deal with a few images at the time, matching the production rate of satellite SAR images under the on-demand-and-at-a-cost environment. In 2014 and 2015, improvements to the processing have been implemented that can now deal with the large data flow that is produced by Sentinel-1. Some of the improvements concern increased efficiency or accuracy, some concern fundamentally different approaches.

Then, the report presents the results of long-term (many months to over one year) monitoring of three sea areas: the Arctic (with two sub-areas), the Western Indian Ocean, and the Mediterranean Sea (Chapters 3, 4 and 5). This kind of long-term monitoring is made possible precisely thanks to Sentinel-1's routine data production approach, and conversely needs the capabilities for automatic bulk processing of its data discussed in Chapter 2. The Arctic is eminently amenable to satellite surveillance, as the satellite orbits converge there (like meridians) so that Arctic (and Antarctic) enjoy the highest revisit frequency. In addition, the Arctic is receiving particular attention in the current years due to the impacts of climate change that are felt there.

The Mediterranean is of course vital to European interests, and as it is always in the Sentinel-1 acquisition scheme, a huge amount of images are available. In fact, finally only a fraction of all the images that Sentinel-1 has taken over the Mediterranean in 2015 have been included in the analysis, the limitation being set by data transfer and storage capacity.

The Western Indian Ocean was surveyed in relation to the interests of the EU in the piracy problem off Somalia (for the moment largely suppressed) and in the maritime safety around Africa in general.

In that region, *all* the images acquired by Sentinel-1 during more than one year have been included in the analysis.

The next two chapters deal with the analysis of detailed ship signatures in Sentinel-1 images for ship size estimation (Chapter 6) and for ship classification (Chapter 7). This is challenging, because Sentinel-1 acquires nearly all of its images in wide-area modes, that – at 20 or 50 meter resolution – sacrifice details for swath width. This is a sound choice for maritime surveillance, but it leaves only little information content on the detected ships besides their position.

Chapter 8 briefly discusses two other issues that have been analysed with Sentinel-1, namely: the occurrence of range ambiguities that cause false alarms and have to be dealt with, especially again for the bulk automatic processing; and how one of the Sentinel-1 products (TOPS SLC) is able to reveal inconsistencies in the apparent headings of SAR ships signatures.

Chapter 9 finally makes some concluding remarks.

All the times in this report are UTC. All Sentinel-1 images specify the unique image name.

2 The SUMO automatic ship detector

2.1 SUMO

SUMO is JRC's automatic ship detection software. It is a software package that reads a satellite SAR image, and analyses it in order to find ships at sea. Its output is a list of detected maritime targets, with their attributes. It uses a land mask to prevent finding targets on land. It has a user interface to display the results, and to allow an operator to edit the results. That editing mostly consists of removing false alarms – detected targets that the automatic algorithm found, but an operator can recognise to be not ships.

Ships manifest themselves in a SAR image as clusters of bright pixels against a darker but noisy background. The background noise is due to reflections from the sea surface ("clutter") and due to radar wave interference effects ("speckle"). These two combined may result in bright pixel clusters, unrelated to ships, that can produce false alarms. Also features at sea that cause a radar echo but that are not ships, such as breaking waves, protruding rocks, reefs, etc., may result in false alarms. The automatic ship detection is often not able to distinguish such a false alarm from a real ship. A human operator can do that much better, mostly because he can take into account other information than what is extracted from the local pixel cluster, such as the contextual information from the image or from knowledge of the region. Therefore, fully automatic ship detection results are usually subjected to a subsequent step of checking by an operator, thereby becoming manually verified results that are much more reliable. Nonetheless, there are also false alarms that cannot be recognised as such even by an experienced human operator. Furthermore, SAR images come with vastly different resolutions, between 1 meter for Spotlight mode and 100 m or more for extra-wide-area modes. The signature of a target can be just a blob at low resolution, as opposed to an extended area with much internal structure at high resolution. There is not yet an automatic ship detection algorithm that works well across the whole range of available resolutions.

SAR images can have multiple polarization channels. For Sentinel-1, images are single polarisation (one channel) or dual polarisation (two channels). In the latter case, each channel is analysed separately by SUMO, and the results then are combined. The SUMO implementation is such that a target in a dual polarization image will be detected if it is has a radar reflection in either one of the channels, and its final extent is the union of the extent in the two channels.

The main attributes of a detected target include position (latitude, longitude), signal strength, size (length, width) and heading (with 180 degree ambiguity).

The DECLIMS project has provided an overview of the state-of-the-art of SAR ship detectors in 2004-2006. The capabilities and performance of the ship detectors on the SAR images at that time, and the shortfalls, were described, including detection rates, false alarm rates, main false alarm causes, and accuracy of size estimation. Much of the DECLIMS conclusions are still valid today, even if the project was carried out some time ago [G2004, G2006]. SUMO was described in [A2010], but since then it has developed further.

2.2 After SUMO – The Blue Hub

For maritime surveillance, the SUMO output (a list of ships detected in a certain satellite SAR image, with their attributes) is not the final product. The SUMO output is ingested by the Blue Hub, the JRC's maritime surveillance development platform [BH2015], for further processing.

SUMO itself works on one image, not taking into account other images. However, the SUMO results (detected targets and their attributes) are uploaded into the Blue Hub database. From there, some further actions are taken. First, SAR images are often delivered as a series of images that are acquired one after the other along the satellite track; a target that falls on the image edge can be detected twice in consecutive images. When consecutive images are stored in the database, such twice detected targets are checked for and removed.

Secondly, the detected targets can be compared with ship positions that are known from other sources, mainly from ship self-reporting data (AIS, LRIT and VMS). This process is referred to as correlation. The ship position reports from AIS, LRIT and VMS arrive only at certain times, so the position of the known ships at the time of the satellite image acquisition must be computed by interpolation (or extrapolation). Another aspect that has to be taken into account in the correlation is that in a SAR image, a moving target is displaced; this is an artefact of the SAR imaging process. Any radial velocity component of the target (i.e., movement towards or away from the radar) leads to a displacement of the target signature in the azimuth direction (i.e., parallel to the radar track). Typical ship speeds (tens of knots) lead to significant displacements (hundred of meters). The exact amount of displacement depends on radar image parameters (such as incidence angle) and ship speed and direction. Therefore, the ship's position and its velocity vector known from the ship reporting data, together with the image parameters, have to be converted to the apparent position of the ship in the image as part of the correlation process.

The outcome of the above processing steps, done in the Blue Hub, is a set of ships with their attributes, including the qualification whether they are seen only by the reporting systems, only by the SAR, or by both. In particular, the non-reporting ships, i.e. the ones that are detected by the SAR but not present in the reporting systems, are of interest.

SUMO does not do ship classification, i.e., it does not include the ship type (like fishing, cargo, tanker, …) among its attributes. The only step toward classification is the ship's size, which is provided. And even size estimation is already difficult due to the nature of the SAR image. Nonetheless, attempts have been performed to do classification on SAR target signatures, based on the SUMO outputs plus the original image data. This is further discussed in Chapter 7.

2.3 SUMO developments

SUMO is continuously being developed at JRC since around 2000. Its first implementation was in IDL [K2001], and after that several new implementations in java have followed. Several algorithms have been developed using matlab, and when mature, been translated in to java.

Over the years, as new satellite SARs have become available, software to ingest their images has been added (SAR image providers inevitably offer images from a new satellite in a new format). Hence, SUMO can now read images from the satellite SARs Radarsat-1, Radarsat-2, ERS-2, ENVISAT-ASAR, Sentinel-1, Cosmo-Skymed, TerraSAR-X and ALOS-PALSAR-2. The facility to read Sentinel-1 was added in 2014, and ALOS-PALSAR-2 in 2015.

SAR images are offered by the providers in various levels or products. The image product can have detected pixels (only amplitude) or complex pixels (with an amplitude and a phase). The phase can be used to measure distances or changes in distance with accuracies of a fraction of a wavelength (which for Sentinel-1 is 5.7 cm). For maritime applications, as the sea level is constantly moving on scales much larger than 5 cm, this is mostly not helpful. There are some maritime scenarios where phase information is useful, but they need multiple radar antennas which Sentinel-1 does not use. The use of phase information with a single antenna for ship detection is still the subject of research, but to date, no convincing results have been produced. Therefore, SUMO uses detected SAR data and discriminates targets only on the basis of higher pixel values than their surroundings. For many satellites, SUMO can ingest the complex image product, but then it converts the pixels into detected values as the first step.

A main cause of false alarms is inaccuracy of the land mask (= coastline). For a long time, the GSHHS global coastline database was used [W1996, W2015]. In 2015, the possibility to use also the OpenStreetMap (OSM) coastline was added [OSM2015]. In many areas of the world, OSM seems to better fit the SAR images than GSHHS; but not everywhere.

To avoid many false alarms, it is often necessary to buffer the land mask outward, i.e., to extend the land mask for some hundreds of meters. This masks out small protrusions of the coast or small islands and even rocks that are not included in the original coast line, but that give rise to false

alarms. Up to recently, this buffering was done on the fly, with each image that was being analysed. In 2015, several pre-buffered land masks were created and stored for faster availability. A buffered land mask also has the advantage that it is coarser, i.e. it needs less points so is faster to load and process.

Apart from the choice between detected or complex, SAR image products can be provided in different projections. The natural projection is on range and azimuth which are the intrinsic SAR image formation coordinates. Further choices include: ground projected, where the pixel frame is not (range, azimuth) but instead a map coordinate with North up such as (longitude, latitude) or (UTM-x, UTM-y); or terrain-geocoded, where the image is corrected for terrain height effects. For ship detection, the preferred product is natural (range, azimuth) projection, because there are no terrain height effects and the natural projection uses no further model assumptions or re-gridding. In that case, information is needed to link the (pixel, line) position of a pixel in the image to a corresponding (longitude, latitude) location on earth. This information must be provided with the image in the metadata. For the Sentinel-1 image product, this is done in the form of a 2D grid of pixels for which the longitude, latitude is specified; the geolocation grid. In 2014, it was noted that the Sentinel-1 geolocation grids often did not have enough accuracy [GS2014]. Ship detection as such does not place very high demands on accuracy. There is no need to obtain the geographic location of a ship to much better than a few tens of meters, because a ship is already often bigger than that; and a ship typically moves, so that its position will have changed by much more than a few tens or hundreds of meters by the time the ship detection results are available. However, the application of an accurate coast line does require a geolocation accuracy of the order of tens of meters. Therefore, in 2015 an alternative method for the geolocation of the Sentinel-1 images was developed. This method is based on using the orbit data. The Sentinel-1 SAR image and its metadata provide very accurate range (distance to satellite) and zero-Doppler (azimuth location) of each pixel , and very accurate 3-D orbit position of the satellite as a function of time. This information allows to calculate the geographic location of each pixels as desired to a high accuracy, higher than provided by the geolocation grid. This method was implemented in SUMO. For now, it is specific to Sentinel-1, and is not implemented in SUMO for other satellites, although it could be extended to them. Figure 2-1 compares the application of land masking to a Sentinel-1 image using the geolocation grid and the orbit data, in both cases using the OpenStreetMap coastline with a buffer of 250 meters. The increased accuracy of the second method is very noticeable.

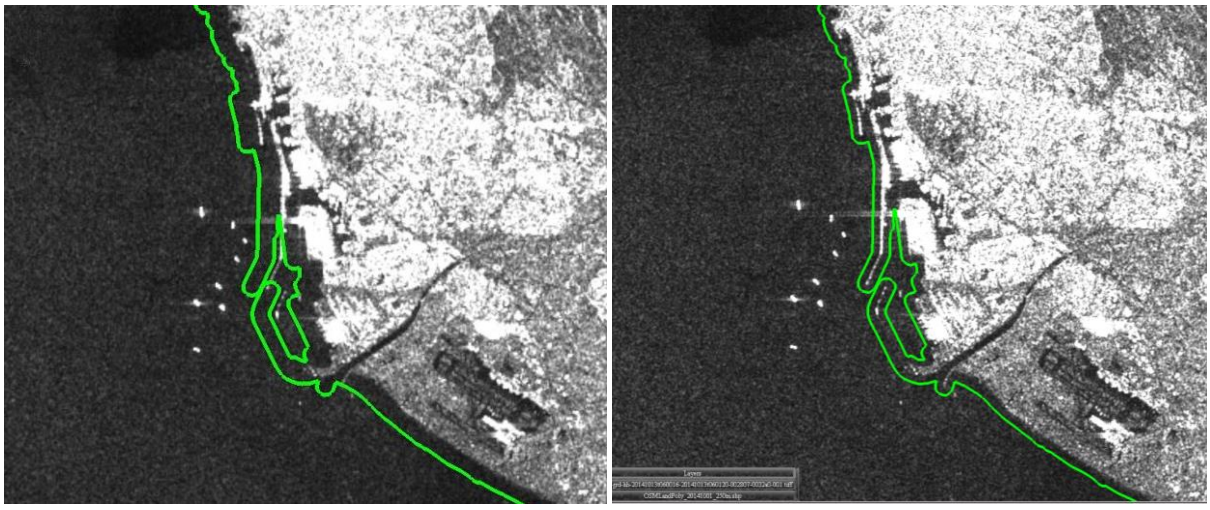

Figure 2-1. Land masking using the geolocation grid (left) and the restituted orbit data (right). The OpenStreetMap coastline with a buffer of 250 meters was used. S1A_EW_GRDH_1SDH_20141013T060016_20141013T060120_002807_0032A0_F436

A particular artefact of SAR images are so-called ambiguities. This is the repeat (re-occurrence) of a target, at a lower level of intensity and at a determined distance from the original. Range ambiguities are the re-appearances of a target at fixed distances in range direction; azimuth ambiguities are the re-appearances of a target at fixed distances in azimuth direction. The distance for range ambiguities is of the order of hundred kilometres. Therefore, they can quickly fall outside of the image (which has a typical size of several tens to several hundreds of km), and conversely, the cause of a range ambiguity that falls inside an image often lies outside the image. The distance for azimuth ambiguities is of the order of 5 km, so they usually can be seen to accompany a bright target against a dark background as two spots roughly 5 km up and 5 km down in azimuth direction. The intensity of an ambiguity can be of the order of 10-100 times lower than the original, so the ambiguities of targets in a relatively bright background are usually not distinguishable. However, for ship detection, especially strong reflectors on land can cause azimuth ambiguities near the coast (on the darker sea background); and they can cause range ambiguities that are however out at sea much further from land. The ambiguity distances are determined by the parameters of the SAR. SUMO has a functionality implemented that flags suspected azimuth ambiguities by checking if there is a brighter target at the known azimuth ambiguity distance from a detection. It has no such functionality for range ambiguities, because their causes often fall outside the image, so cannot be verified.

SUMO was updated in 2015 to include the flagging of azimuth ambiguities for Sentinel-1. Moreover, it had been found in 2014 that Sentinel-1 Interferometric Wide (IW) images have additional ghosts in azimuth to the classical azimuth ambiguity [GS2014]. These are manifested as bright and often doubly-peaked spots away in azimuth from a bright target, but closer than the classical azimuth ambiguity distance. Their cause was unknown. In 2015, the pattern of their occurrence was analysed, and it was determined that they occur at a constant distance for a given sub-swath and a given Sentinel SAR processor version, as shown in Table 2-1. Their flagging was implemented in SUMO, so that SUMO results now suppress the detection of false alarms caused by both the classical azimuth ambiguity as well as the additional ghost in azimuth of Sentinel-1 images. In the mean time, ESA has communicated that these new azimuth ghosts are thought to be the result of some approximation in the SAR processing, and that a new release of the processing software should alleviate this problem in 2016.

Table 2-1. Estimated azimuth distance (in meters) of Sentinel-1 IW's double peak ambiguity to originating target, as a function of sub-swath and Sentinel SAR processor (IPF) version.

	IPF versions up to and including v236	**IPF versions after v236**
IW1	3,817	2,417
IW2	4,310	2,951
IW3	4,162	2,624

Most Sentinel-1 images have a dark border of invalid pixels whose values are low but not necessarily 0 (usually in range from 0 to 10). If included in the ship detection analysis they cause many false detections, but they cannot be easily excluded from the analysis since their pixel values are not much lower than the values of sea pixels. Since the image metadata do not indicate the location of the non-valid pixels (at least in the non-complex products), an algorithm based on adaptive thresholding has been implemented in SUMO in 2015 to remove these borders. This algorithm: (a) scans the areas near the four borders of the image; and (b) in each of the four areas it finds the cut-off points that separate non-valid pixels from valid-pixels. Using the cut-off values, the pixels in the dark area are excluded from the analysis. The algorithm has proved very successful, as shown in Figure 2-2, which compares the ship detection result when the dark border is not excluded from the

analysis with when it is excluded. When it is not excluded, the analysis results in many detected pixels (green dots in the figure); when it is excluded, the false detections are no longer there.

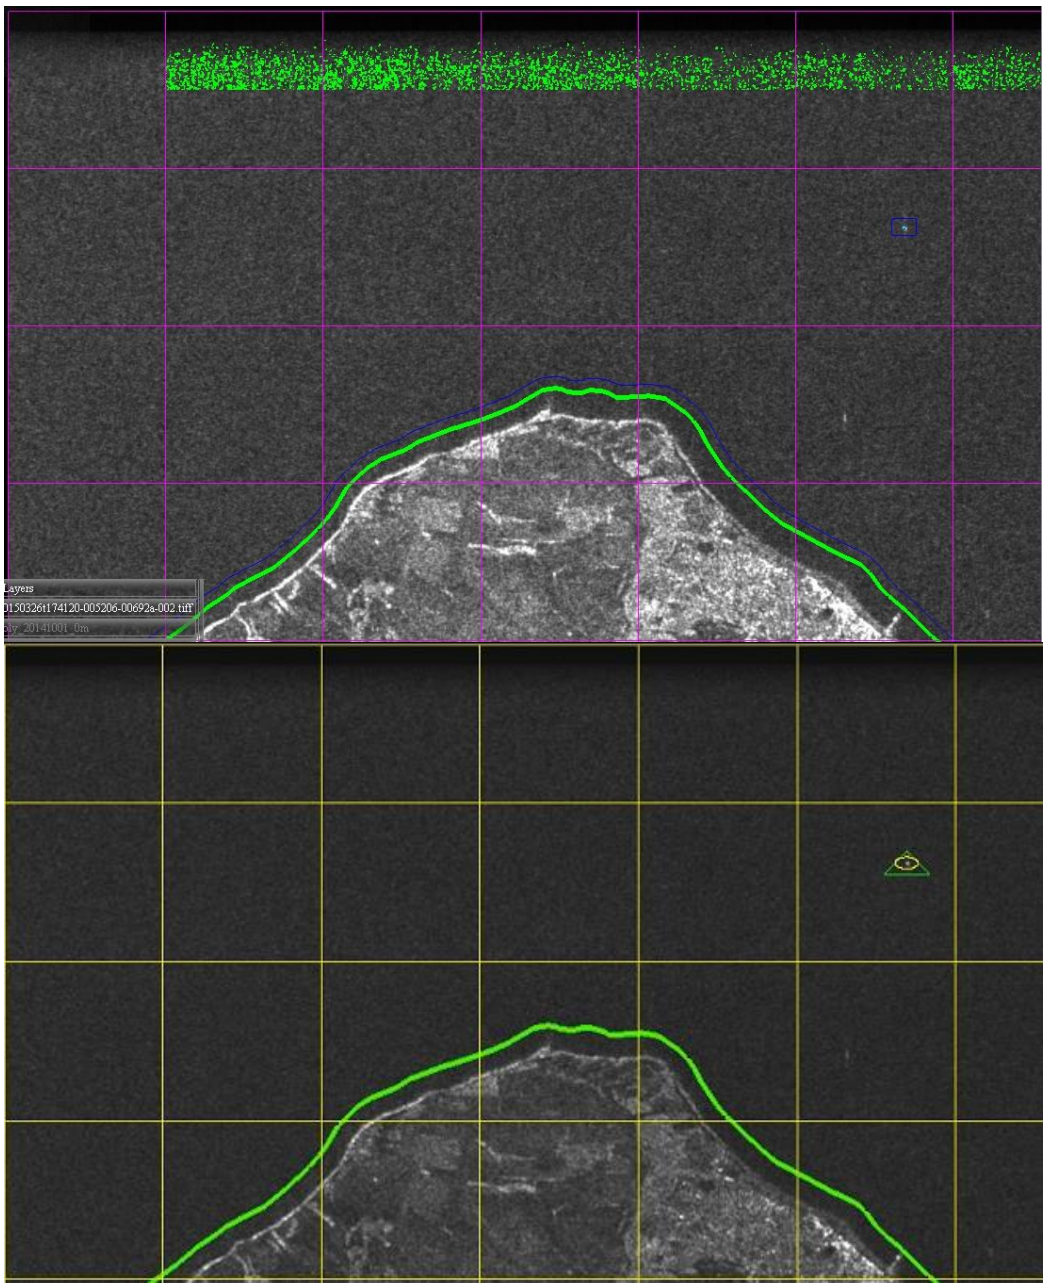

Figure 2-2. Results of the ship detection analysis when the dark border around a Sentinel-1 image is not excluded (top) and when it is excluded (bottom). Many false detections (green dots) are returned when the border is not excluded. The grid in the images is just a visual reference. S1A_IW_GRDH_1SDV_20150326T174051_20150326T174120_005206_00692A_F72A

Also implemented in 2015 is a new framework to identify recurrent targets. These are targets that repeat in the same position over different satellite images. They may be ghost targets or actual fixed structures (like oil platforms), and they are generally of little interest as far as maritime surveillance is concerned. Furthermore, it is better to remove them from the list of SAR detections, otherwise they may give rise to incorrect correlations with reporting ships. This framework can also be used to discriminate between ghost targets and fixed structures, since ghost targets only appear at a given position for a specific set of sensor and geometric parameters (mode, beam, orbit track), whereas fixed structures appear for any set of parameters. Sentinel-1's repeat acquisition operations have

made it possible to test this framework in real campaigns. Further details are given in [SG2015]. This framework is still in development, but it has already been applied in the Western Indian Ocean and Mediterranean Sea campaigns reported in this document.

3 Sentinel-1 in the Arctic

This chapter shows monitoring results of Sentinel-1 over two locations in the Arctic, revealing how shipping patterns are linked to sea ice conditions.

3.1 Yenisei Gulf

This is a gulf on the Kara Sea. It is the estuary of the Yenisei River that cuts into Siberia.

From the Sentinel-1 rolling archive, 67 images over the Yenisei Gulf were selected that covered the red box in Figure 3-1. The images are all EW mode and HH polarisation, but mixed GRDH and GRDM products. They were acquired between 9 October 2014 and 20 April 2015. The times between the consecutive images vary between 1.4 and 9.0 days, with an average of 2.9 days.

From all images, the area of the red box was cut out, and this was smoothed to fixed resolution and resampled in a North-up grid. Thereby, the GRDH and GRDM products resulted in 67 similar images.

Figure 3-1. Yenisei Gulf in Siberia, with the outline of the Sentinel-1 images that were used. The outline size is 120 km by 67 km, and the geographic coordinates of the corners are 72.1N, 77.81632E to 72.7N, 81.33452E (map from Google Earth).

An ice chart of 14 April 2015 confirms that the Gulf was covered in ‘fast ice’ at that day (Figure 3-2).

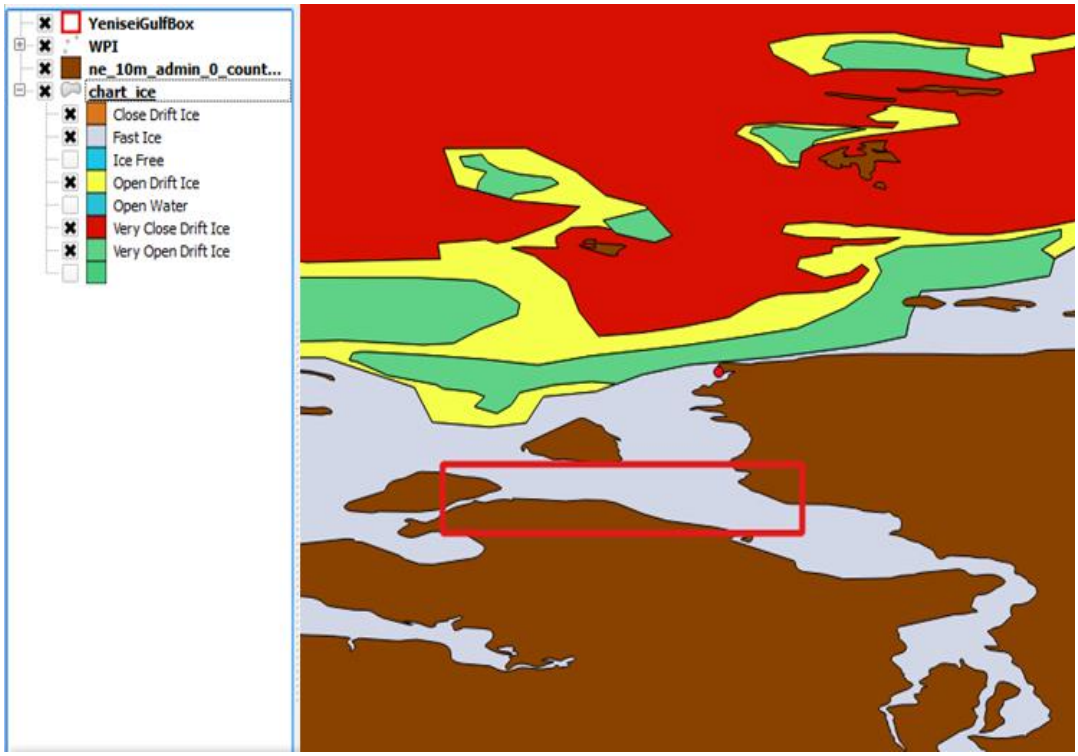

Figure 3-2. Ice chart of 14 April 2015 over the Yenisei Gulf (source: Norwegian Meteorological Institute).

The image series runs throughout the winter, and it shows the appearance and evolution of the winter sea ice. Some examples from the series are shown. The first image, of 9 Oct 2014 (Figure 3-3), still has open water, with no sea ice recognisable. The image of the next day, 10 Oct (Figure 3-4), already shows the first appearance of sea ice. During the following days, the sea ice is characterised by patches that move around. From 19 Oct (Figure 3-5) onward, the ice stays fixed in some locations just on the coast. On 24 Nov (Figure 3-6), a large part of the ice in the middle of the gulf starts to remain fixed. From 6 Dec (Figure 3-7) onward, the ice in the entire gulf stays fixed in place. Also from that image onward, a line starts to appear that can be linked to a ship route crossing the scene, connecting the Yenisei River with the Kara Sea. This line remains recognisable in all the following images, e.g. shown on 5 Mar 2015 (Figure 3-9) and through to the last image of 20 Apr (Figure 3-10). All the while, the sea ice remains frozen in place. At 13 Jan (Figure 3-8), a second line appears that branches off of the first one and becomes even stronger, suggesting that becomes the main traffic route.

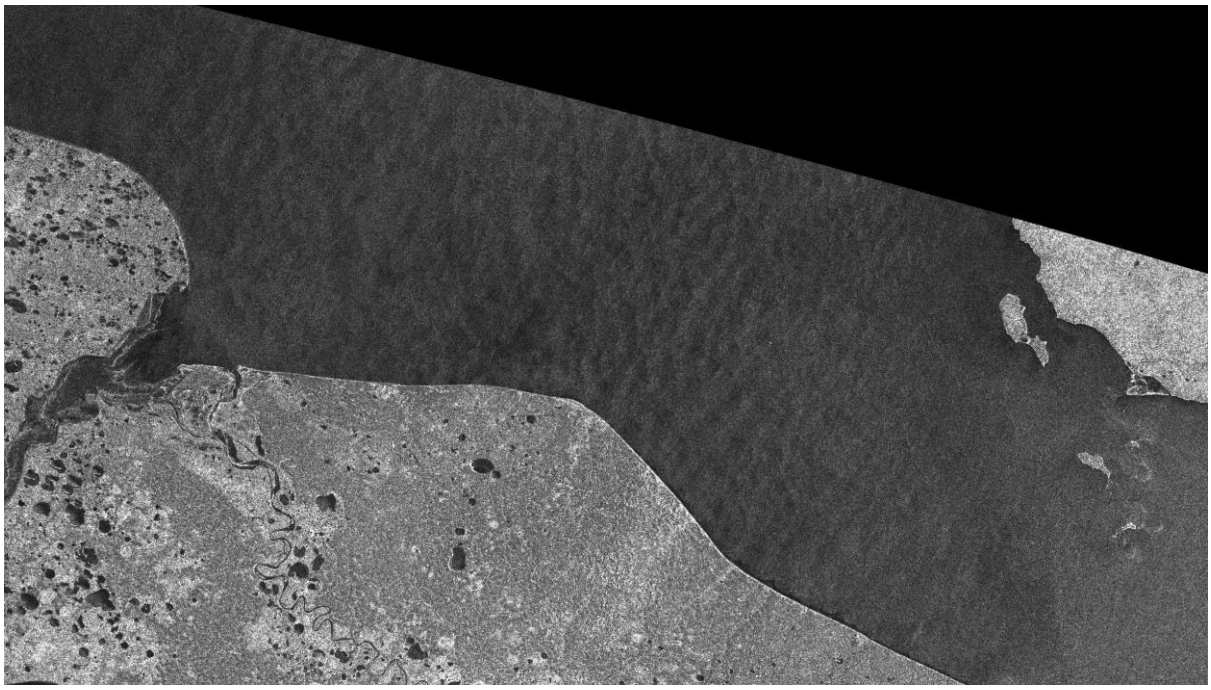

Figure 3-3. Yenisei Gulf, 9 Oct 2014.
S1A_EW_GRDH_1SDH_20141009T012901_20141009T012941_002746_003150_BE72

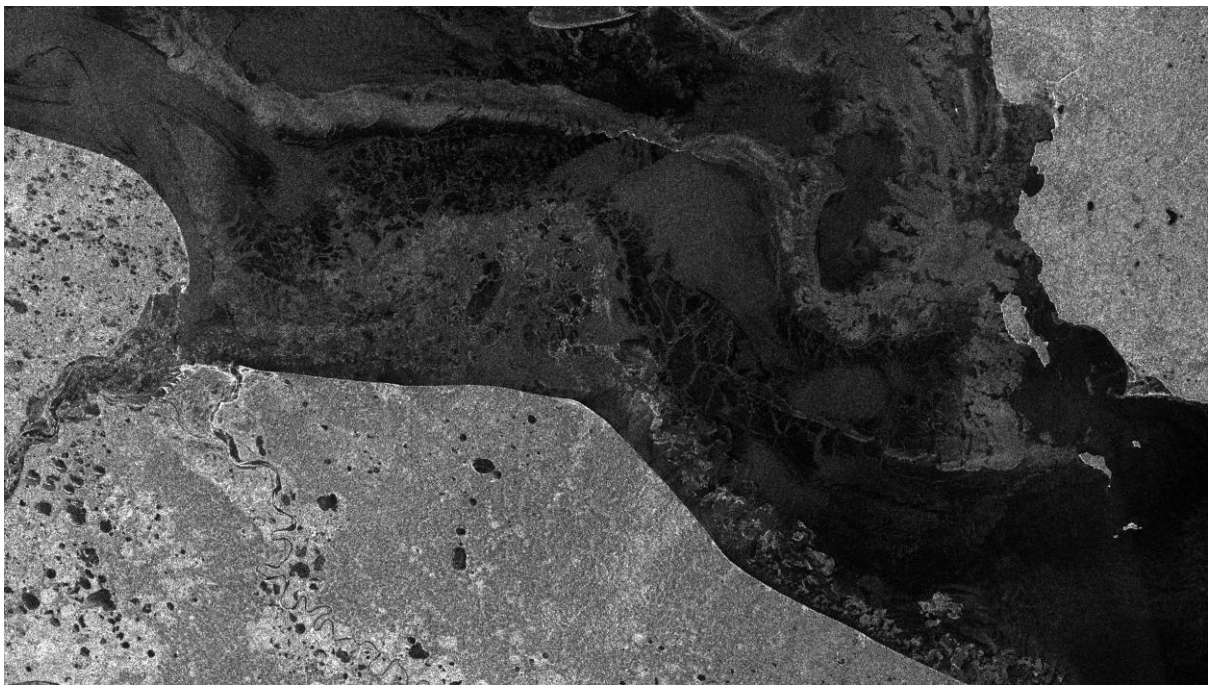

Figure 3-4. Yenisei Gulf, 10 Oct 2014.
S1A_EW_GRDH_1SDH_20141010T115254_20141010T115358_002767_0031C7_38E2

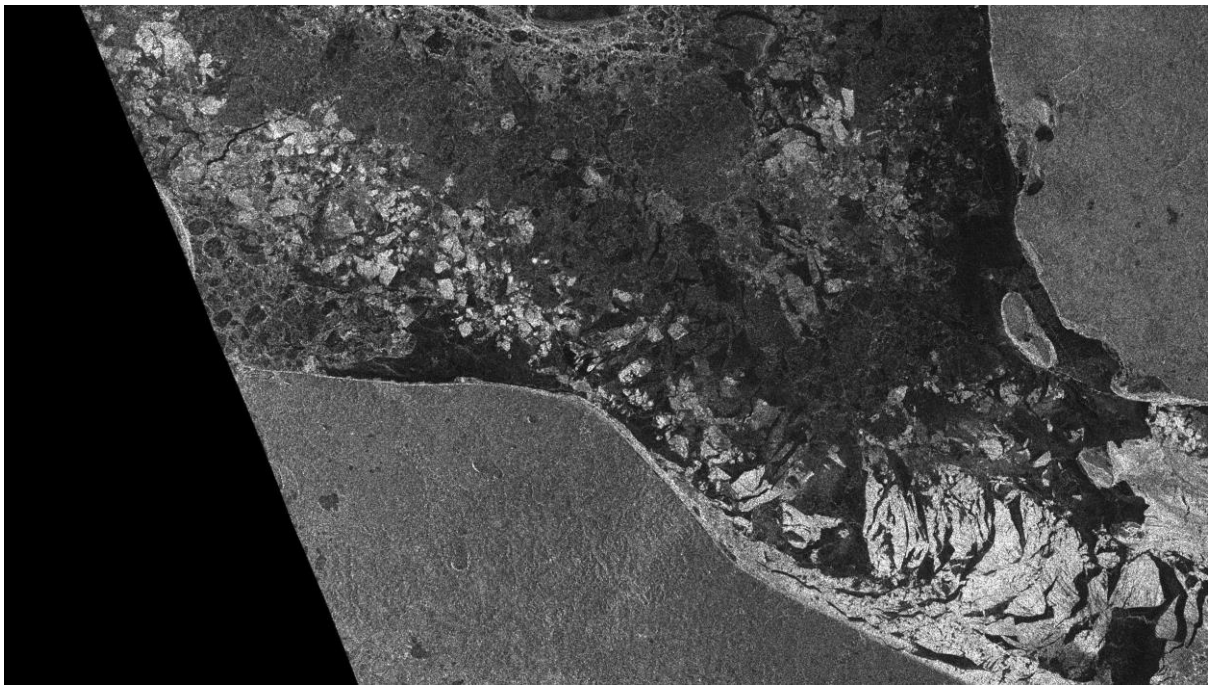

Figure 3-5. Yenisei Gulf, 19 Oct 2014.
S1A_EW_GRDH_1SDH_20141019T112823_20141019T112927_002898_003485_02C3

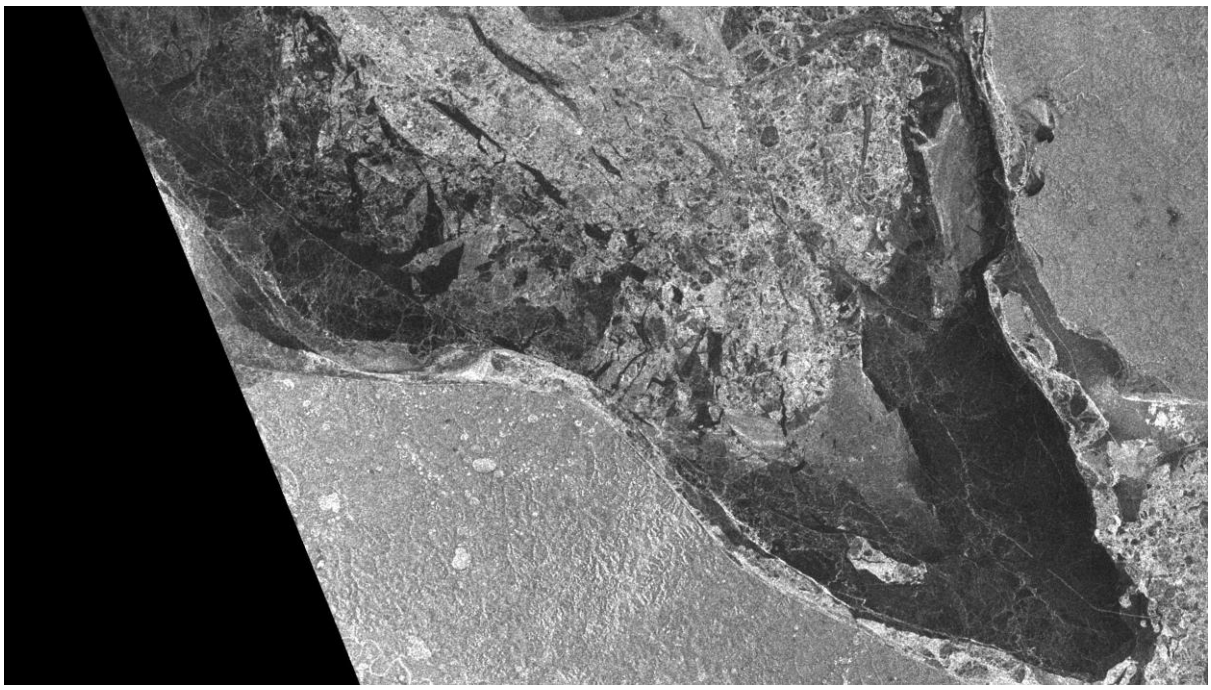

Figure 3-6. Yenisei Gulf, 24 Nov 2014.
S1A_EW_GRDM_1SDH_20141124T112822_20141124T112927_003423_003FEE_962E

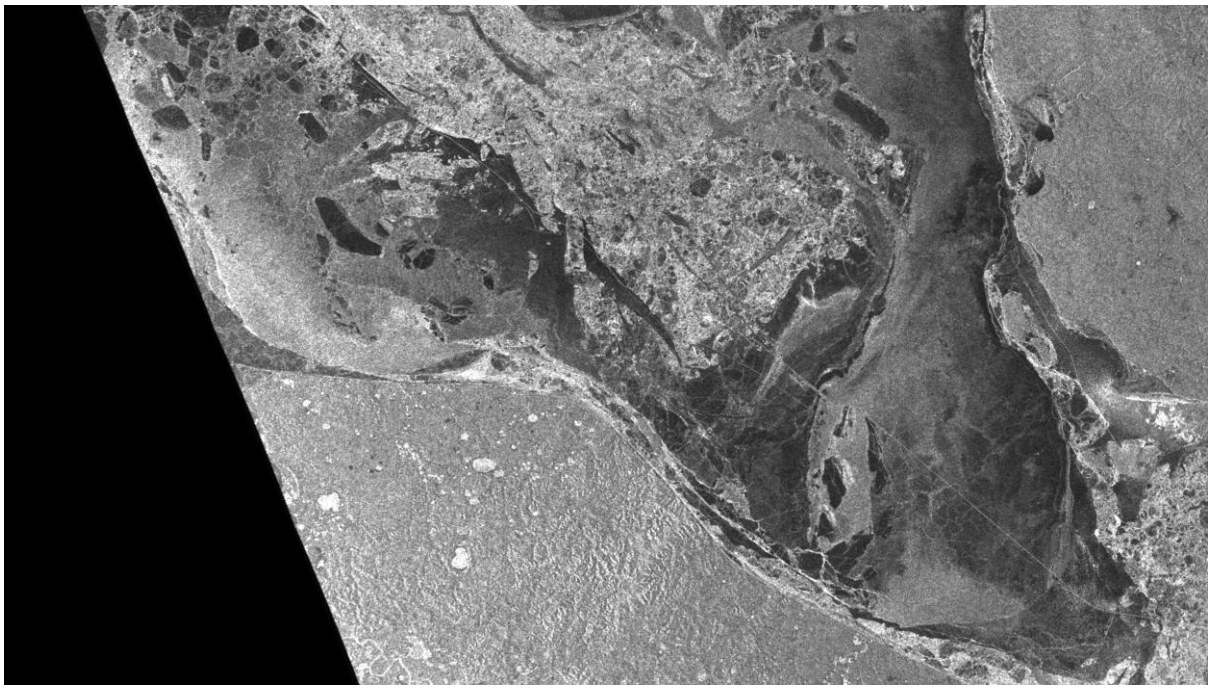

Figure 3-7. Yenisei Gulf, 6 Dec 2014.
S1A_EW_GRDM_1SDH_20141206T112804_20141206T112908_003598_004406_7229

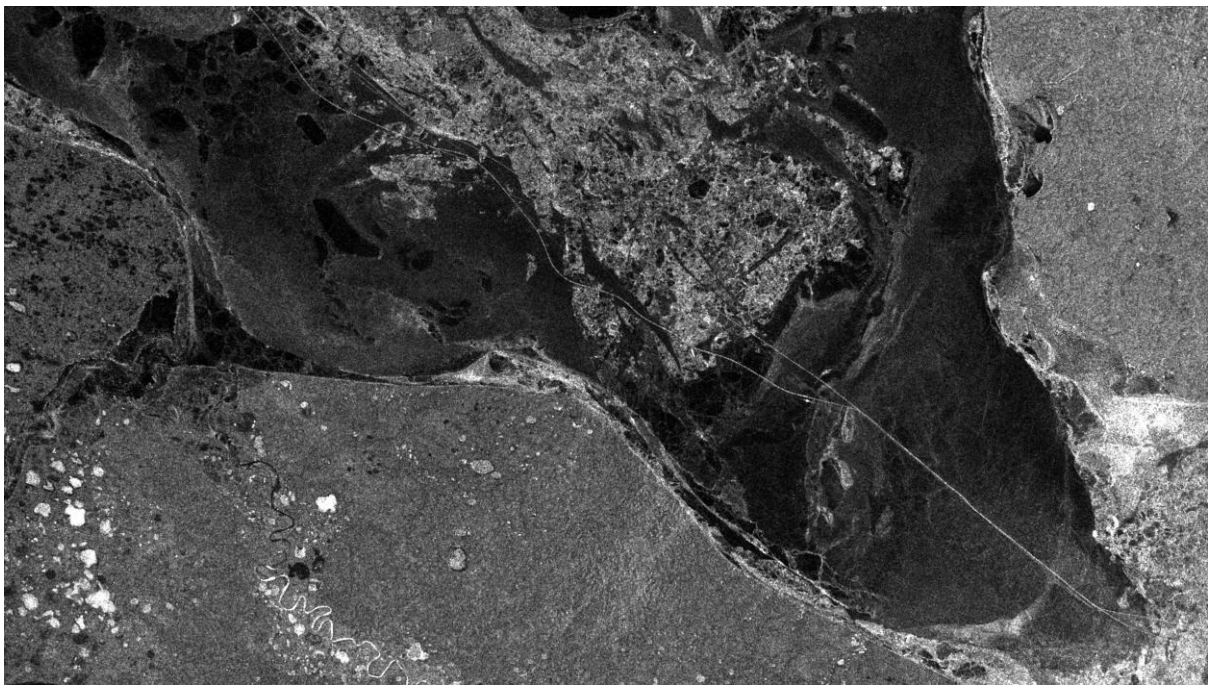

Figure 3-8. Yenisei Gulf, 13 Jan 2015.
S1A_EW_GRDM_1SDH_20150113T012813_20150113T012913_004146_00506B_92E8

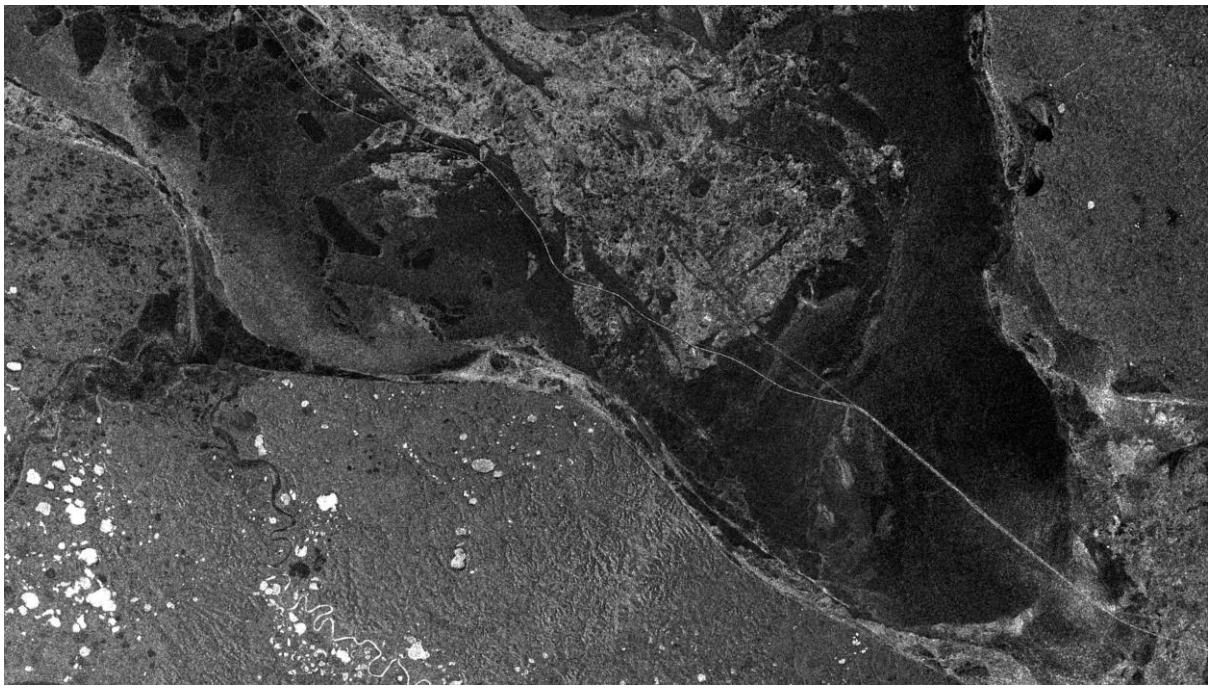

Figure 3-9. Yenisei Gulf, 5 Mar 2015.
S1A_EW_GRDM_1SDH_20150305T113610_20150305T113714_004896_0061AB_4EB8

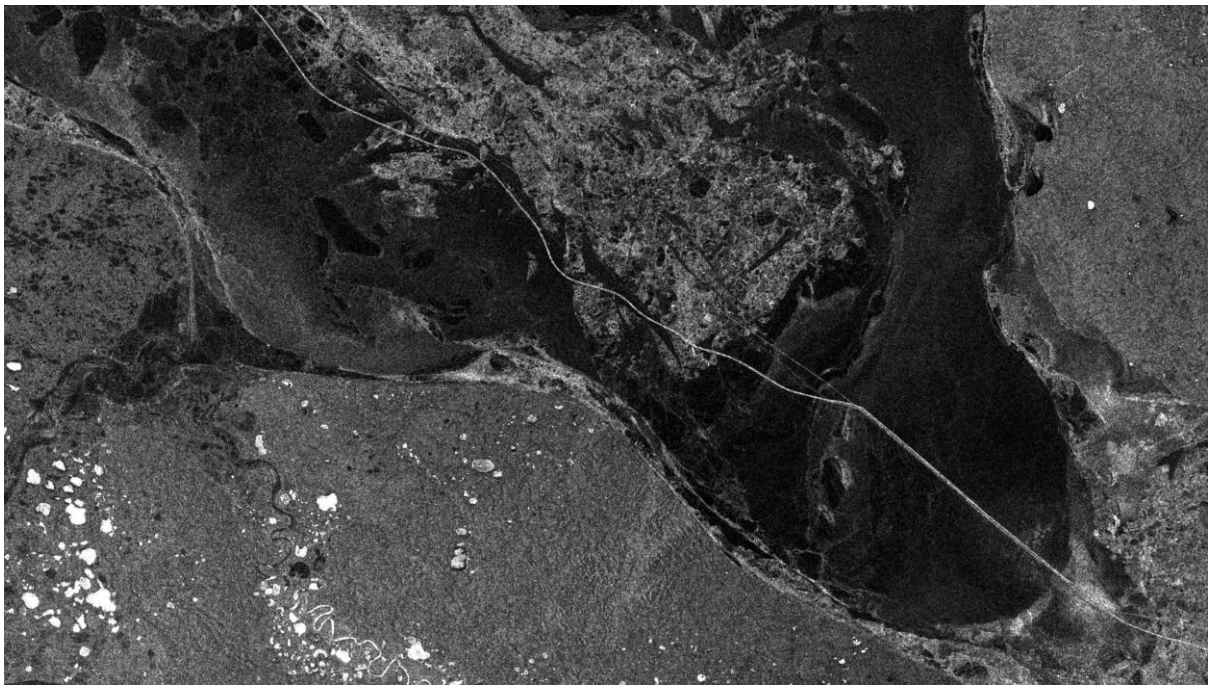

Figure 3-10. Yenisei Gulf, 20 Apr 2015.
S1A_EW_GRDM_1SDH_20150420T115228_20150420T115333_005567_0071F8_CE1F

The reason why the bright line across the Sentinel-1 images is understood to reflect the location of a ship traffic route, is because of AIS messages received. Figure 3-11 plots the locations of more than 5,000 AIS messages that were received during 3.5 months starting 1 Jan 2015 when the ice was frozen in place. With the AIS messages, the ships are identified, but that information is not published here. However, at least one of the MMSI numbers belongs to an icebreaker. It can be seen that almost all messages lie on the bright, somewhat crooked line that runs though the Sentinel-1 images. A few AIS messages lie on the straighter line that appeared earlier in the image sequence. Figure 3-12 shows only one week of AIS messages in March, where all messages lie on the crooked track; while Figure 3-13 shows one week of AIS messages from early September 2014, where all

messages lie on the straight track. The interpretation is that the straight track is the shortest route that was followed when the water was ice free and when the ice was still thin; while the crooked track is a track made by icebreakers through relatively thinner parts of the ice, that has remained open and in use for the entire winter.

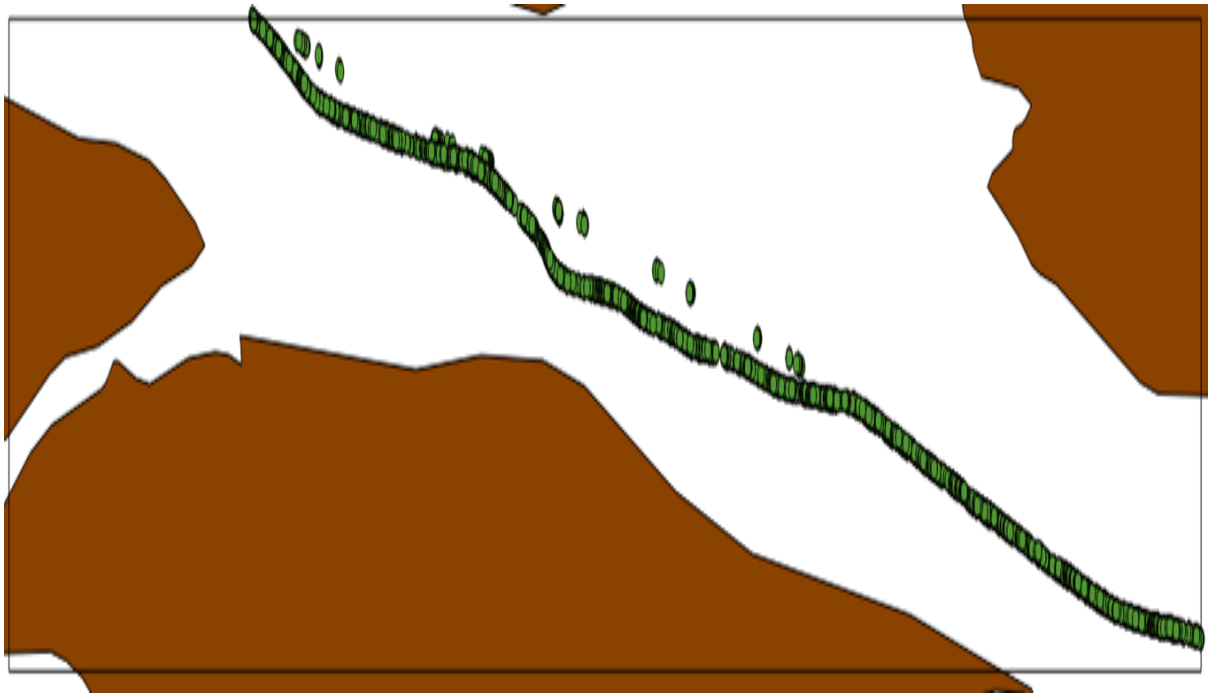

Figure 3-11. AIS messages received between 1 Jan and 22 Apr 2015 (3.5 months) plotted in the same area as the previous Sentinel-1 images. The total number of messages is 5,088. Land is indicated in brown.

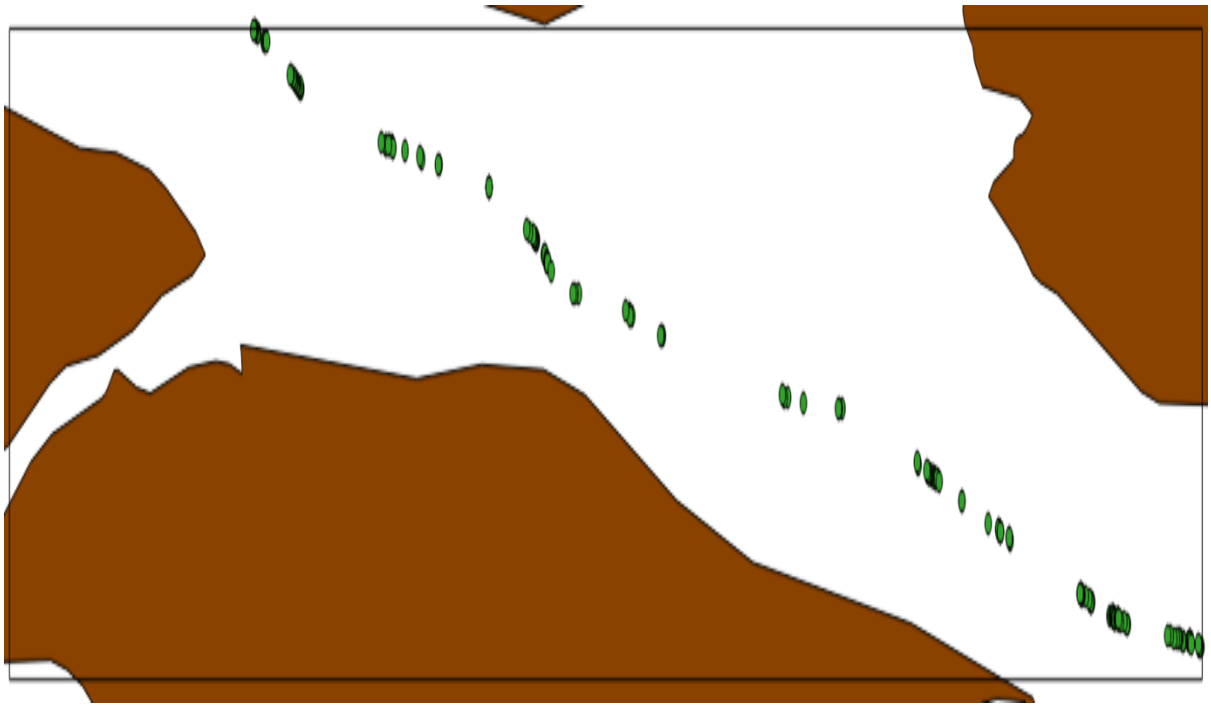

Figure 3-12. A subset of the AIS messages from the previous figure, the ones received during the week 9 – 15 Mar 2015 (177 messages).

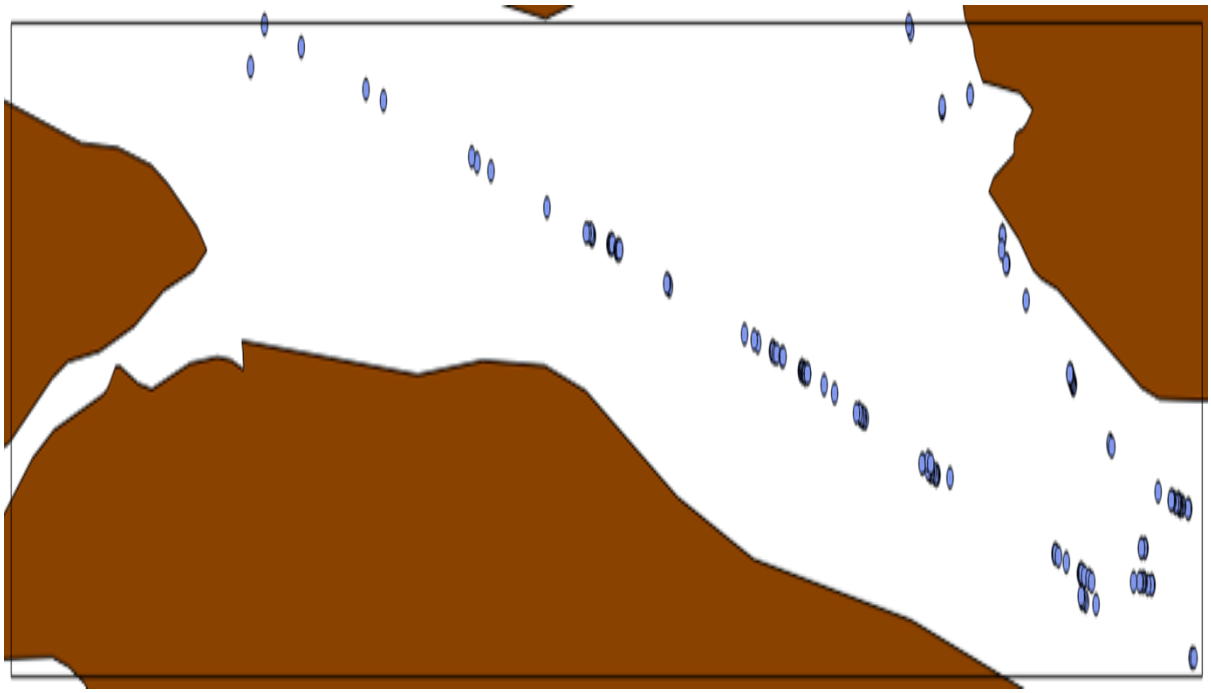

Figure 3-13. A subset of the AIS messages from Figure 3-11, the ones received during the week 1-7 Sep 2014 (157 messages).

Figure 3-14 zooms in on three ships in the Sentinel-1 images. The ships show up as a bright blob, located next to the bright line that is interpreted as the track through the sea ice. This is in accordance with the SAR imaging mechanism, whereby targets that have a speed in range direction (left-right in the images) receive a displacement in the azimuth direction (up-down in the images). Therefore, the moving targets appear to be located off their actual track. The travelling direction of the ships (from the Yenisei river or towards it) can also be determined by whether the targets appear displaced to one side of the track or to the other.

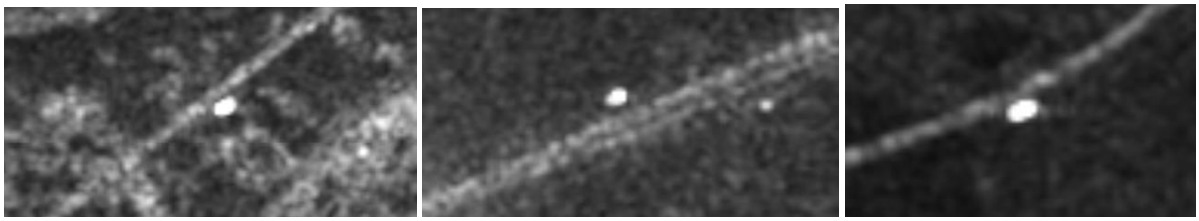

Figure 3-14. Zoom-ins of three ship targets seen next to the track through the sea ice.

3.2 Bay of Bothnia

A similar analysis as above was performed on the Bay of Bothnia, which is the northern part of the Gulf of Bothnia, near the Finnish city of Oulu.

Again, 67 Sentinel-1 images were downloaded, between 1 Jan 2015 and 21 Apr 2015. This covers a total timespan of 110.5 days (3.6 months). The times between the consecutive images vary between 0.5 and 4.5 days, with an average of 1.7 days. Most images (59) are in EW mode / GRDM product / HH polarisation; 8 are in IW mode / GRDH product / VV polarisation. The area corresponds to the red box in Figure 3-15.

An ice chart shows that most of the Bay was covered with ‘fast ice’ on 14 April 2015 (Figure 3-16).

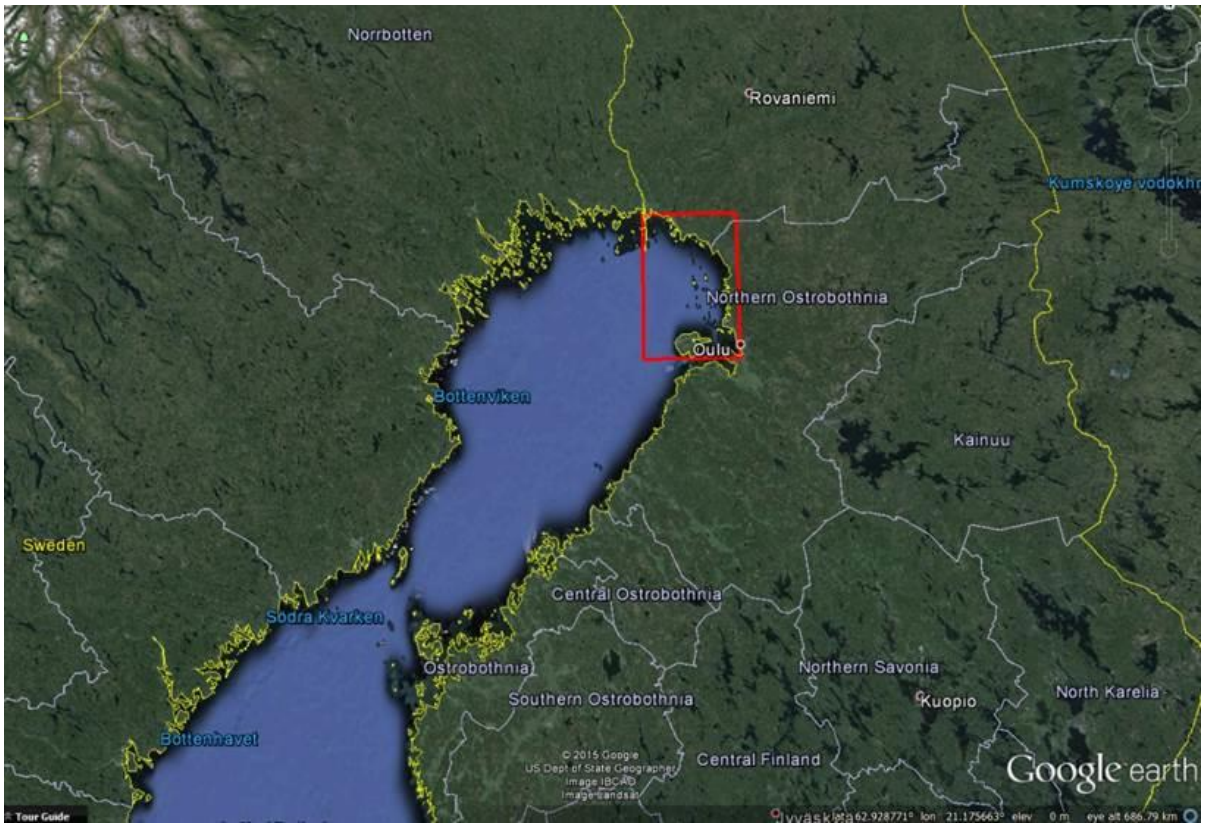

Figure 3-15. Bay of Bothnia, with the outline of the Sentinel-1 images that were used. The outline size is 61 km by 95 km, and the geographic coordinates of the corners are 64.94N, 24.12E to 65.79N, 25.46E (map from Google Earth).

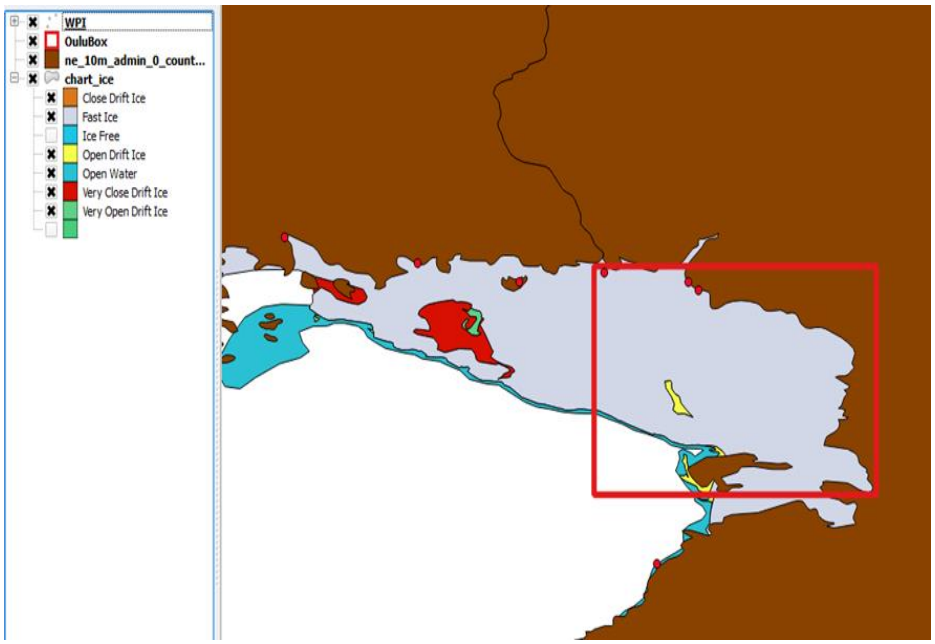

Figure 3-16. Ice chart of the Bay of Bothnia of 14 April 2015, confirming that the part attached to the coast was covered with 'fast ice' that day (source: Norwegian Meteorological Institute).

A few images of the series are shown. The first images from the start of January contain sea ice in many places that remains fixed where it is attached to the coast, but moves or changes where it is further out at sea. Through to mid-February, sea ice away from the coast covers at times large areas, remaining in place for some days but then moving or disappearing again (Figure 3-17 to Figure 3-19).

From the first images onward, bright tracks like the ones seen in the Yenisei Gulf start to form. By mid-January there are many. As before, they are ascribed to shipping activity, because their locations correlate with AIS tracks, as shown in Figure 3-25 and Figure 3-26. One long track runs parallel to the coast, and several point in/out from the coast. In the places where the ice remains fixed, the tracks also remain fixed. In the places where the ice can be seen to move in large chunks, the tracks often remain visible and move with the ice. By mid-February, most of the image area is covered by ice with only some apparently open water in the bottom left corner (Figure 3-20). Nonetheless, parts of the outer reaches of the ice still move back and forth from time to time, the bright tracks remaining coherent (Figure 3-21). Figure 3-22 shows the tracks very clearly. By mid-April, the outer parts of the sea ice break up and disappear, but still ice remains fixed to the coast (Figure 3-23). In the final image of the series (Figure 3-24), the main track parallel to the coast is split into several smaller ones next to each other. Presumably, the ice is now thin enough to move through at will and create new trails.

The following figures are all EW mode and HH polarisation, except for two cases (IW-VV) where it is indicated.

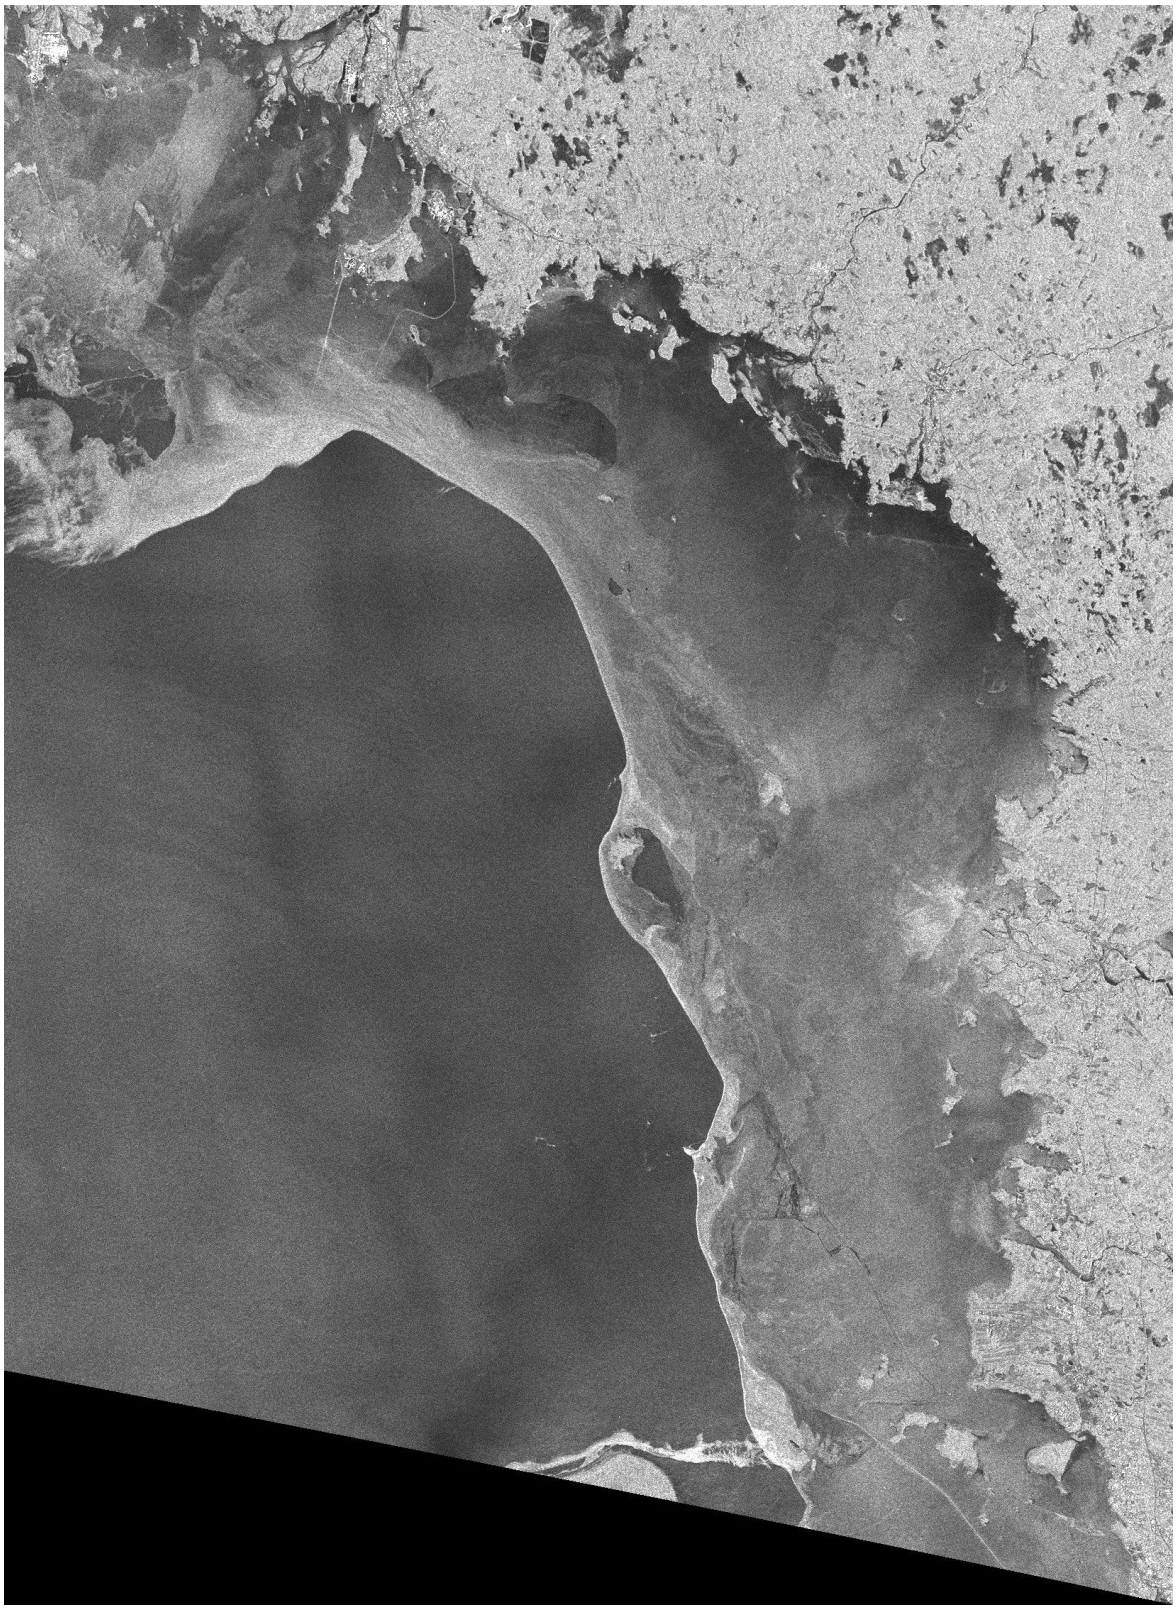

Figure 3-17. Bay of Bothnia, 1 Jan 2015. This image is IW mode VV polarisation unlike most others. The coast of Finland is to the top and to the right. Sea ice is apparently extending out from the coast, with the edge of the ice running though the centre of the image, the bottom/left part being apparently ice-free water.
S1A_IW_GRDH_1SDV_20150101T044812_20150101T044837_003973_004C86_552D

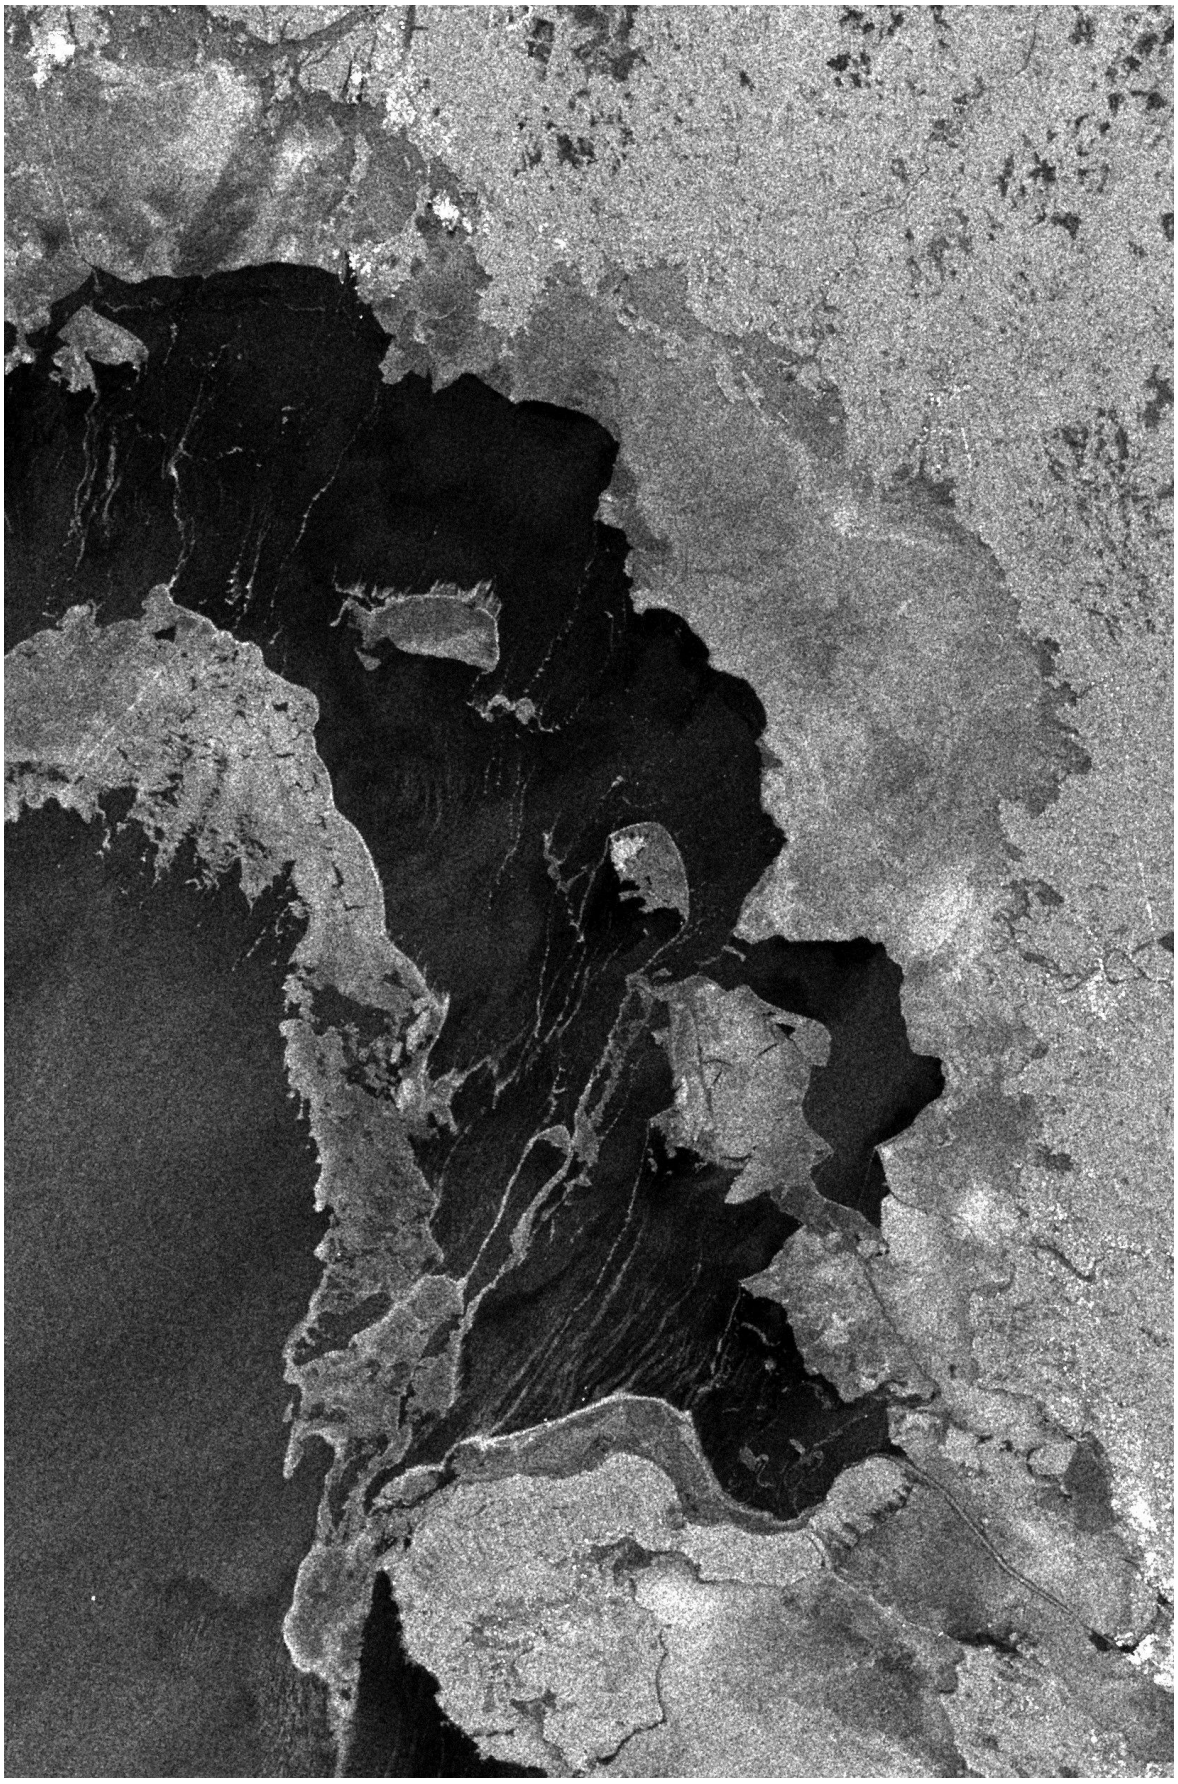

Figure 3-18. Bay of Bothnia, 3 Jan 2015. Compared to two and a half days earlier, large parts of the ice cover seems to have floated out to sea.
S1A_EW_GRDM_1SDH_20150103T154936_20150103T155000_004009_004D44_137B

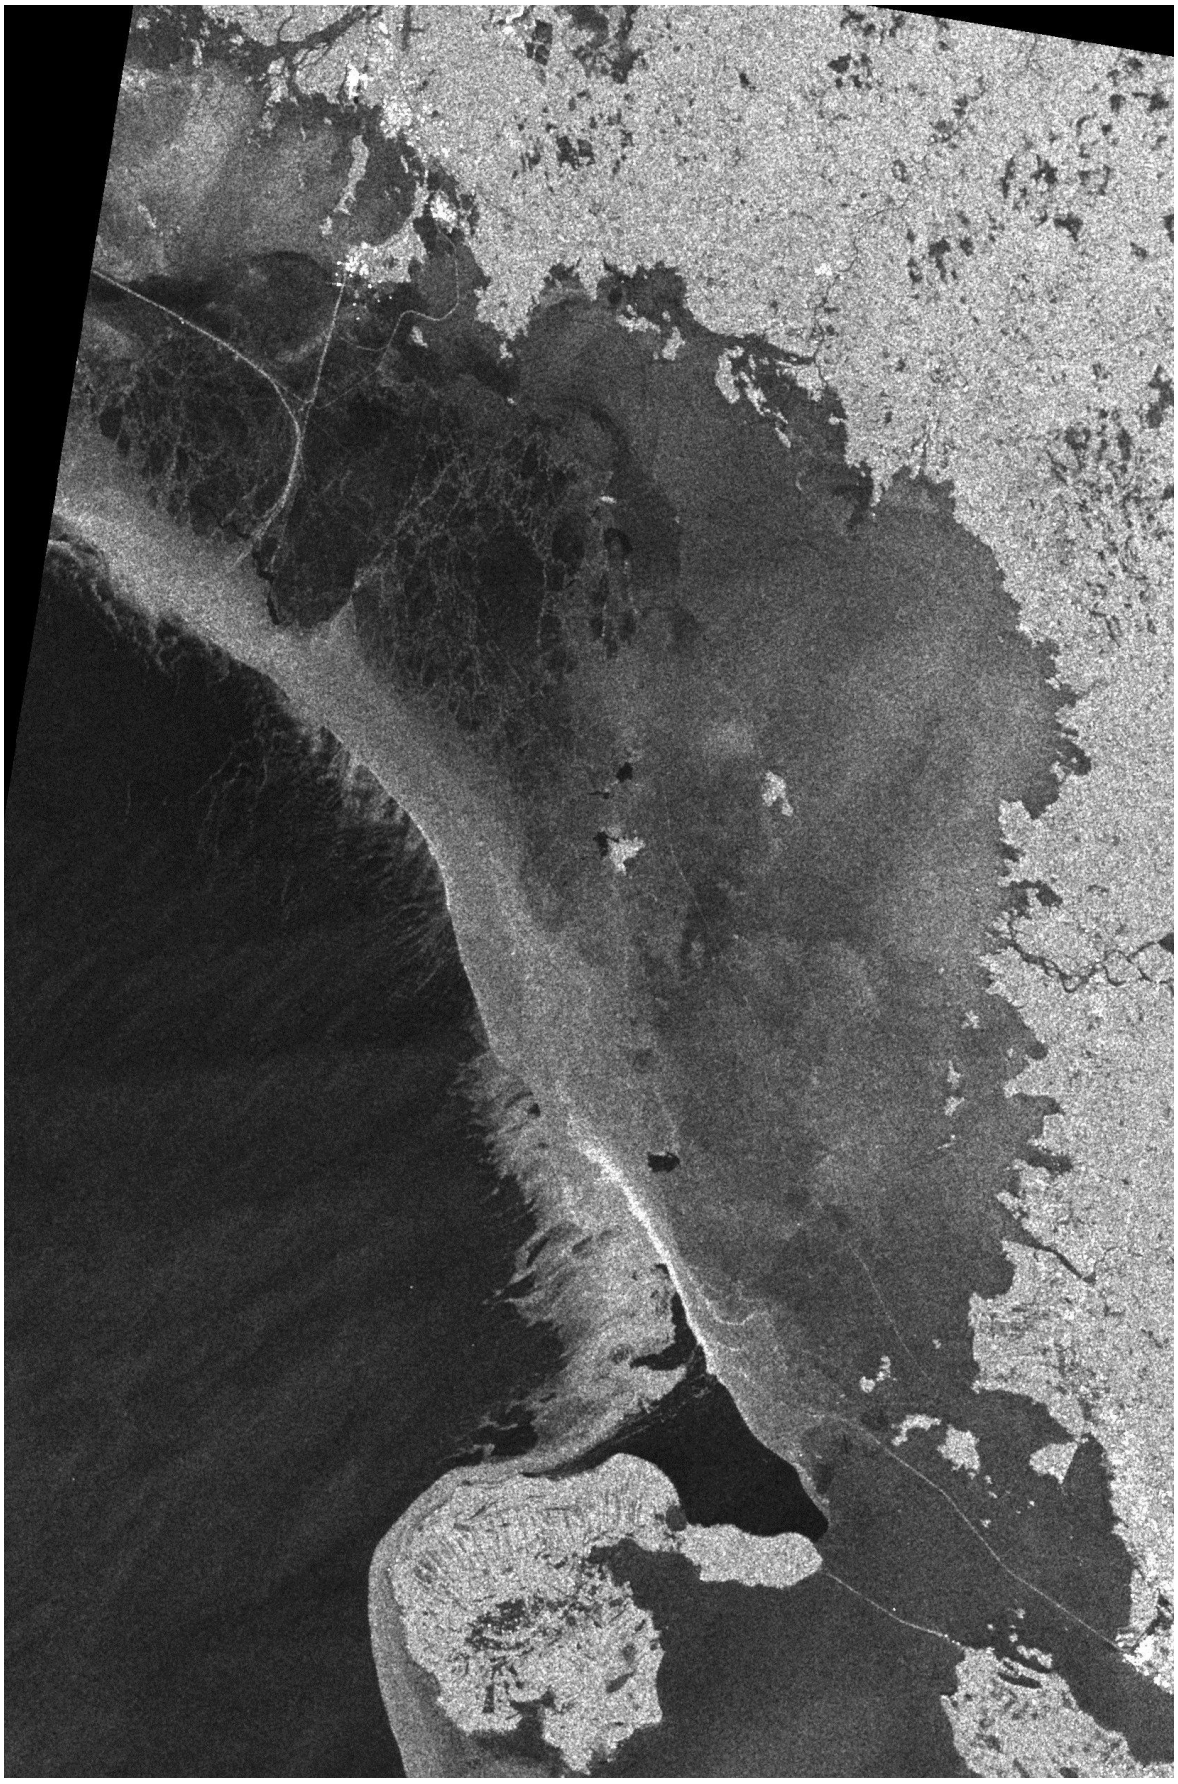

Figure 3-19. Bay of Bothnia, 8 Jan 2015. Bright, thin tracks have started to show up, especially in the top left and bottom right.
S1A_EW_GRDM_1SDH_20150108T044011_20150108T044115_004075_004ECA_67C7

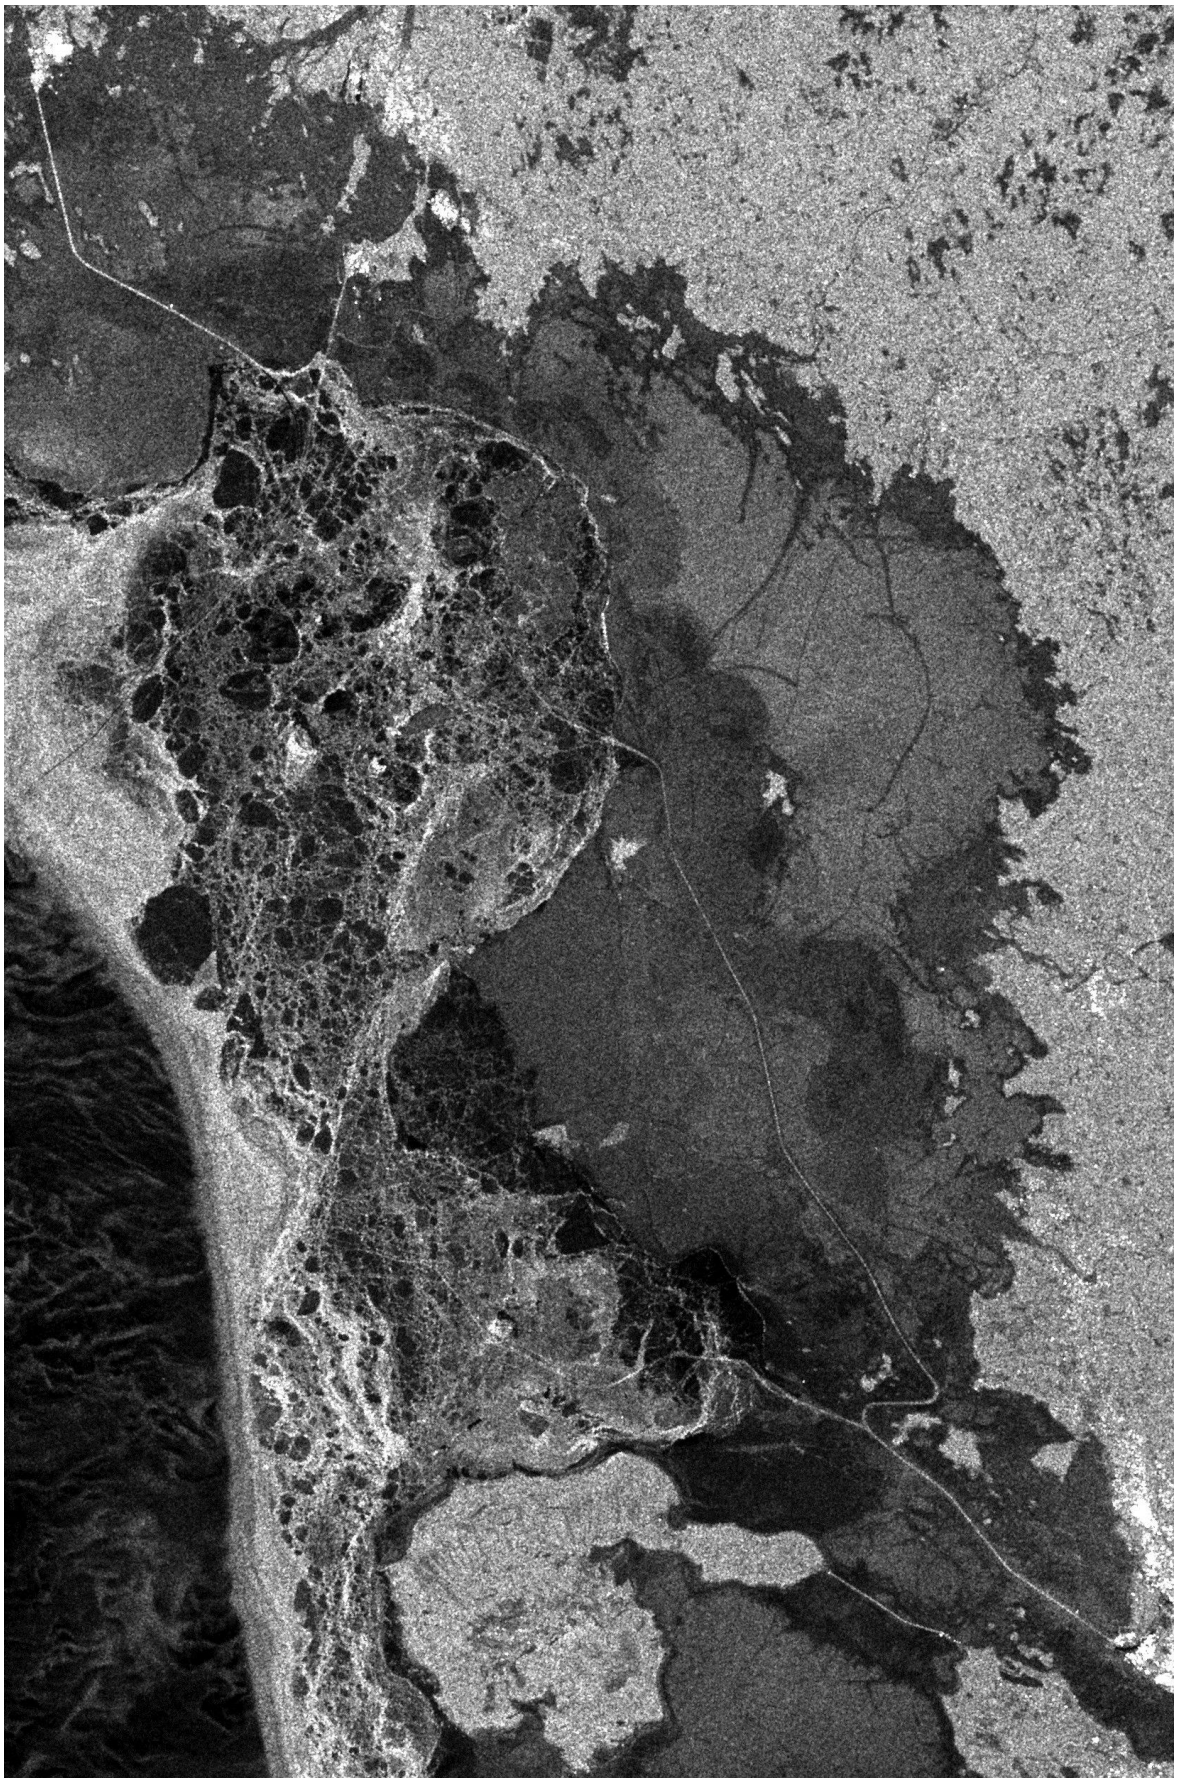

Figure 3-20. Bay of Bothnia, 13 Feb 2015. A long track parallel to the coast has formed. S1A_EW_GRDM_1SDH_20150213T155719_20150213T155814_004607_005ACA_8EB8

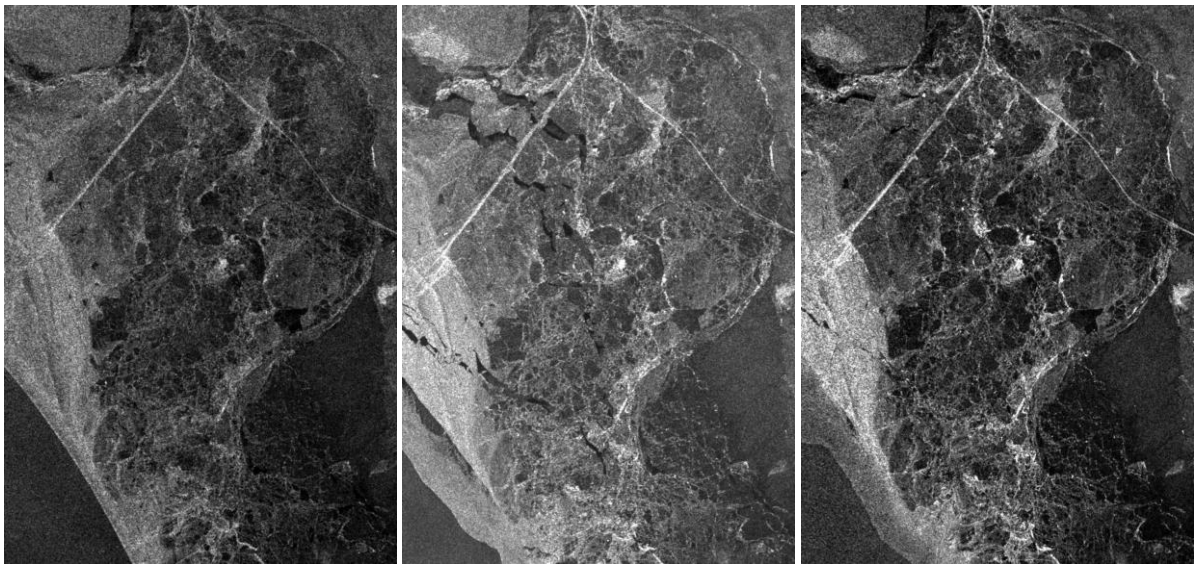

Figure 3-21. Bay of Bothnia, 9, 11 and 12 Mar 2015. This is a part from the centre of the image. On 11 Mar, some cracks have formed, pushing parts of the ice out to sea. But the next day, the cracks are closed again, and the ice structures as well as the bright tracks correlate accurately with those of 3 days before.
S1A_EW_GRDM_1SDH_20150309T155720_20150309T155814_004957_006339_3664
S1A_EW_GRDM_1SDH_20150311T154126_20150311T154151_004986_0063E2_8532
S1A_EW_GRDM_1SDH_20150312T050444_20150312T050548_004994_006421_0876

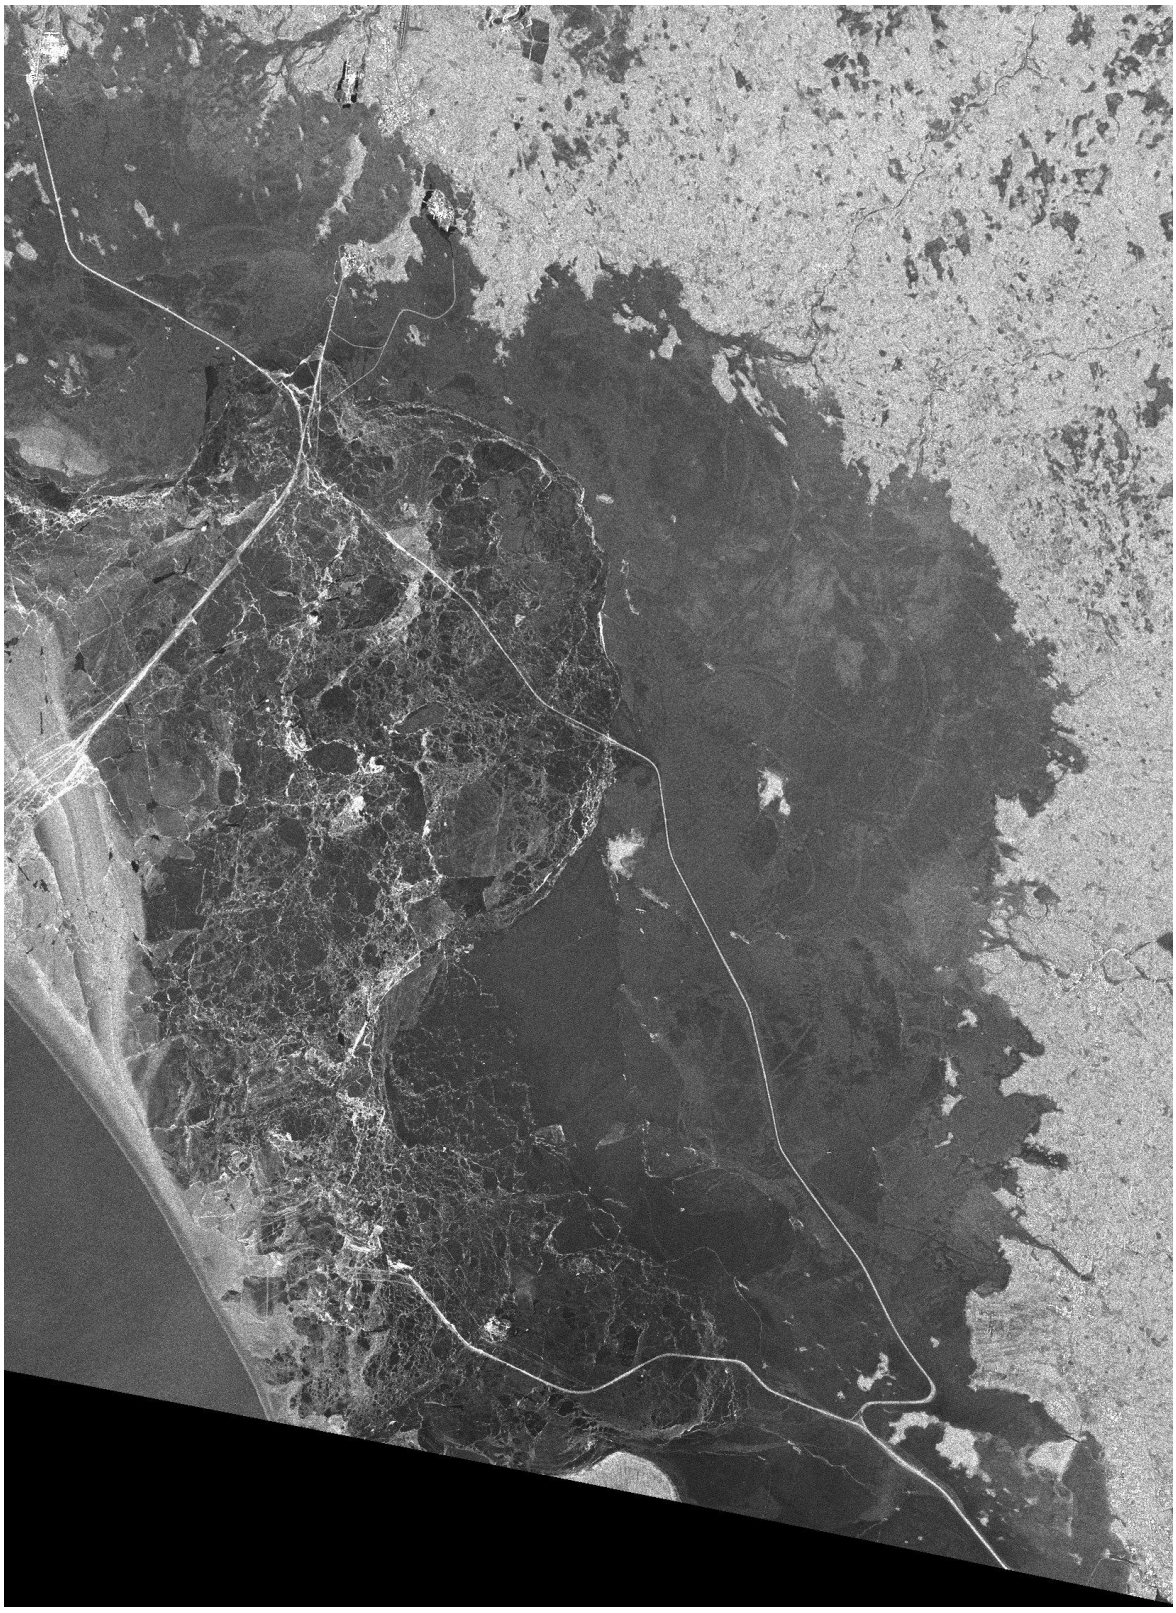

Figure 3-22. Bay of Bothnia, 14 Mar 2015. This image is IW mode VV polarisation unlike most others. The bright tracks are well developed and show up with much contrast. S1A_IW_GRDH_1SDV_20150314T044811_20150314T044836_005023_0064C9_190B

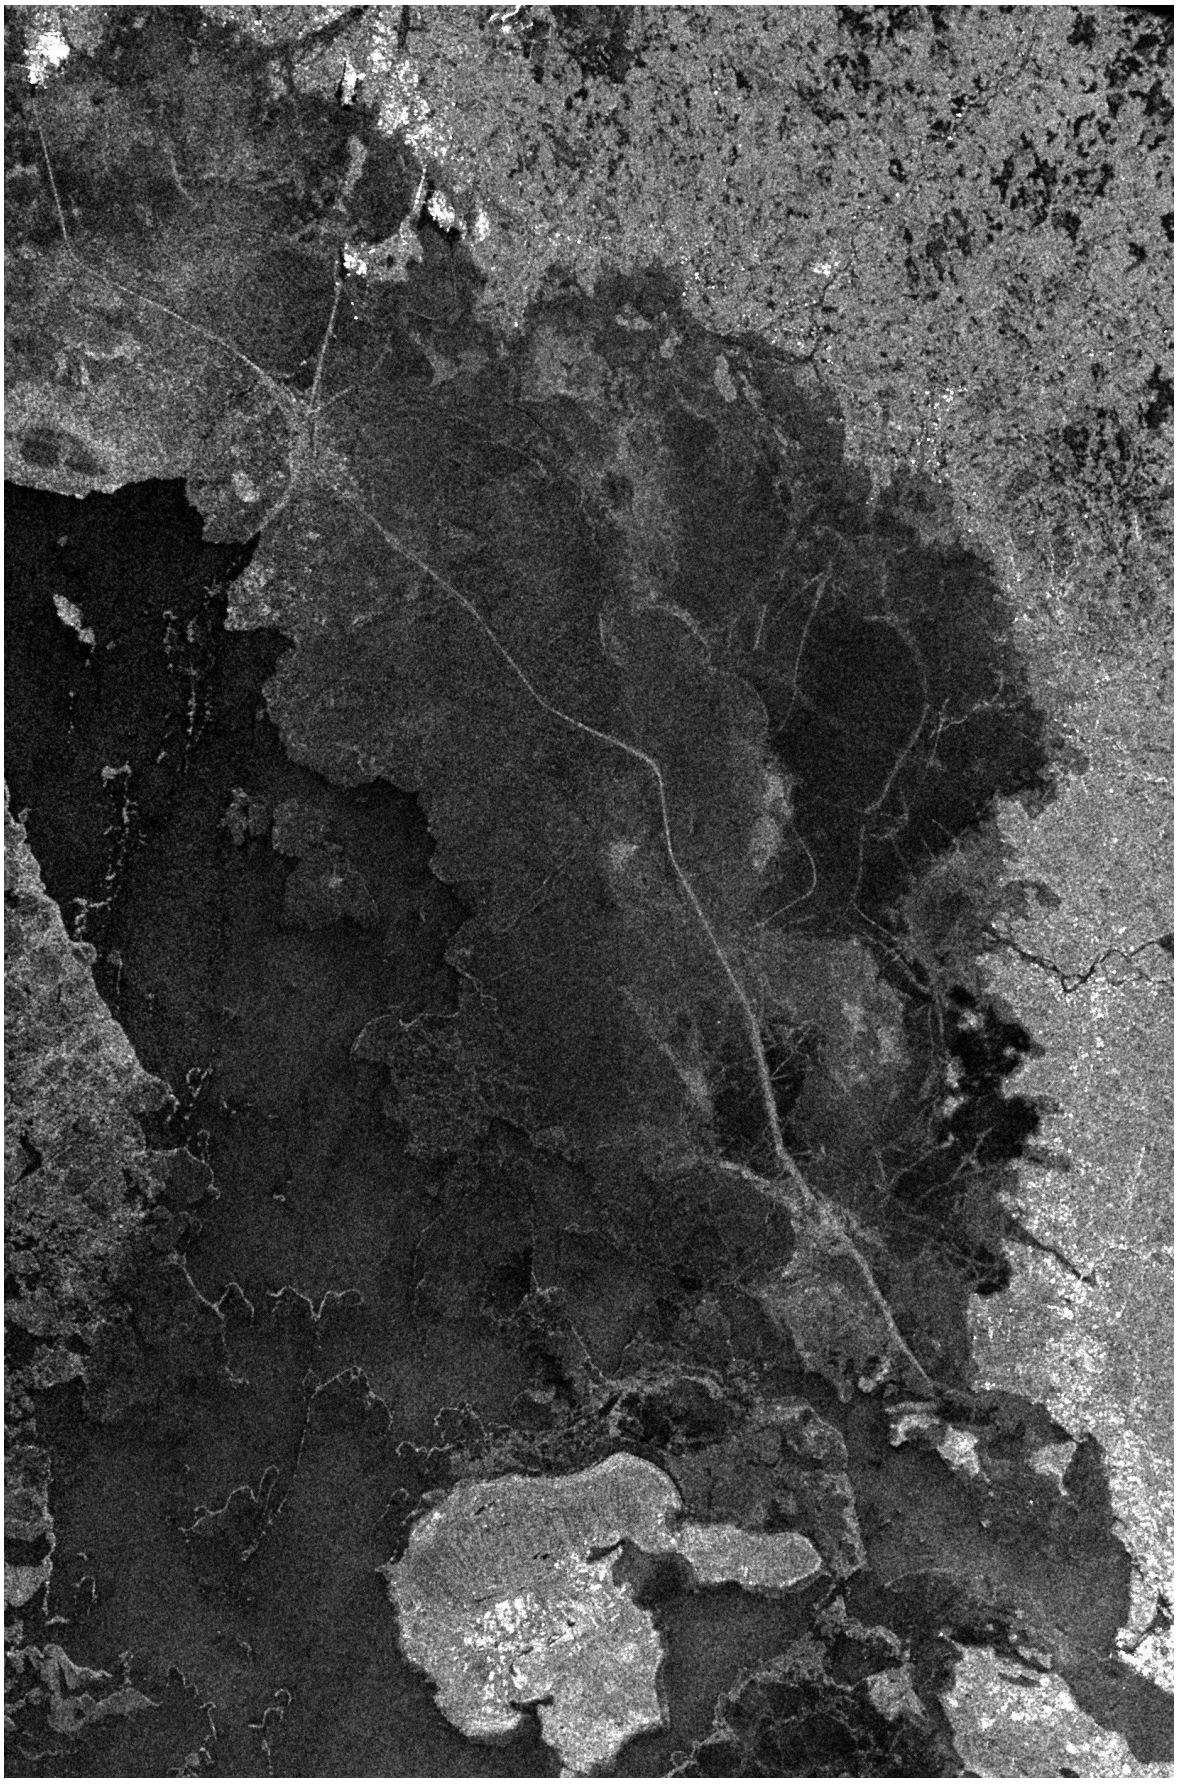

Figure 3-23. Bay of Bothnia, 17 Apr 2015. The ice has started to break up, the outer parts drifting off to sea.
S1A_EW_GRDM_1SDH_20150417T050446_20150417T050550_005519_0070D6_B653

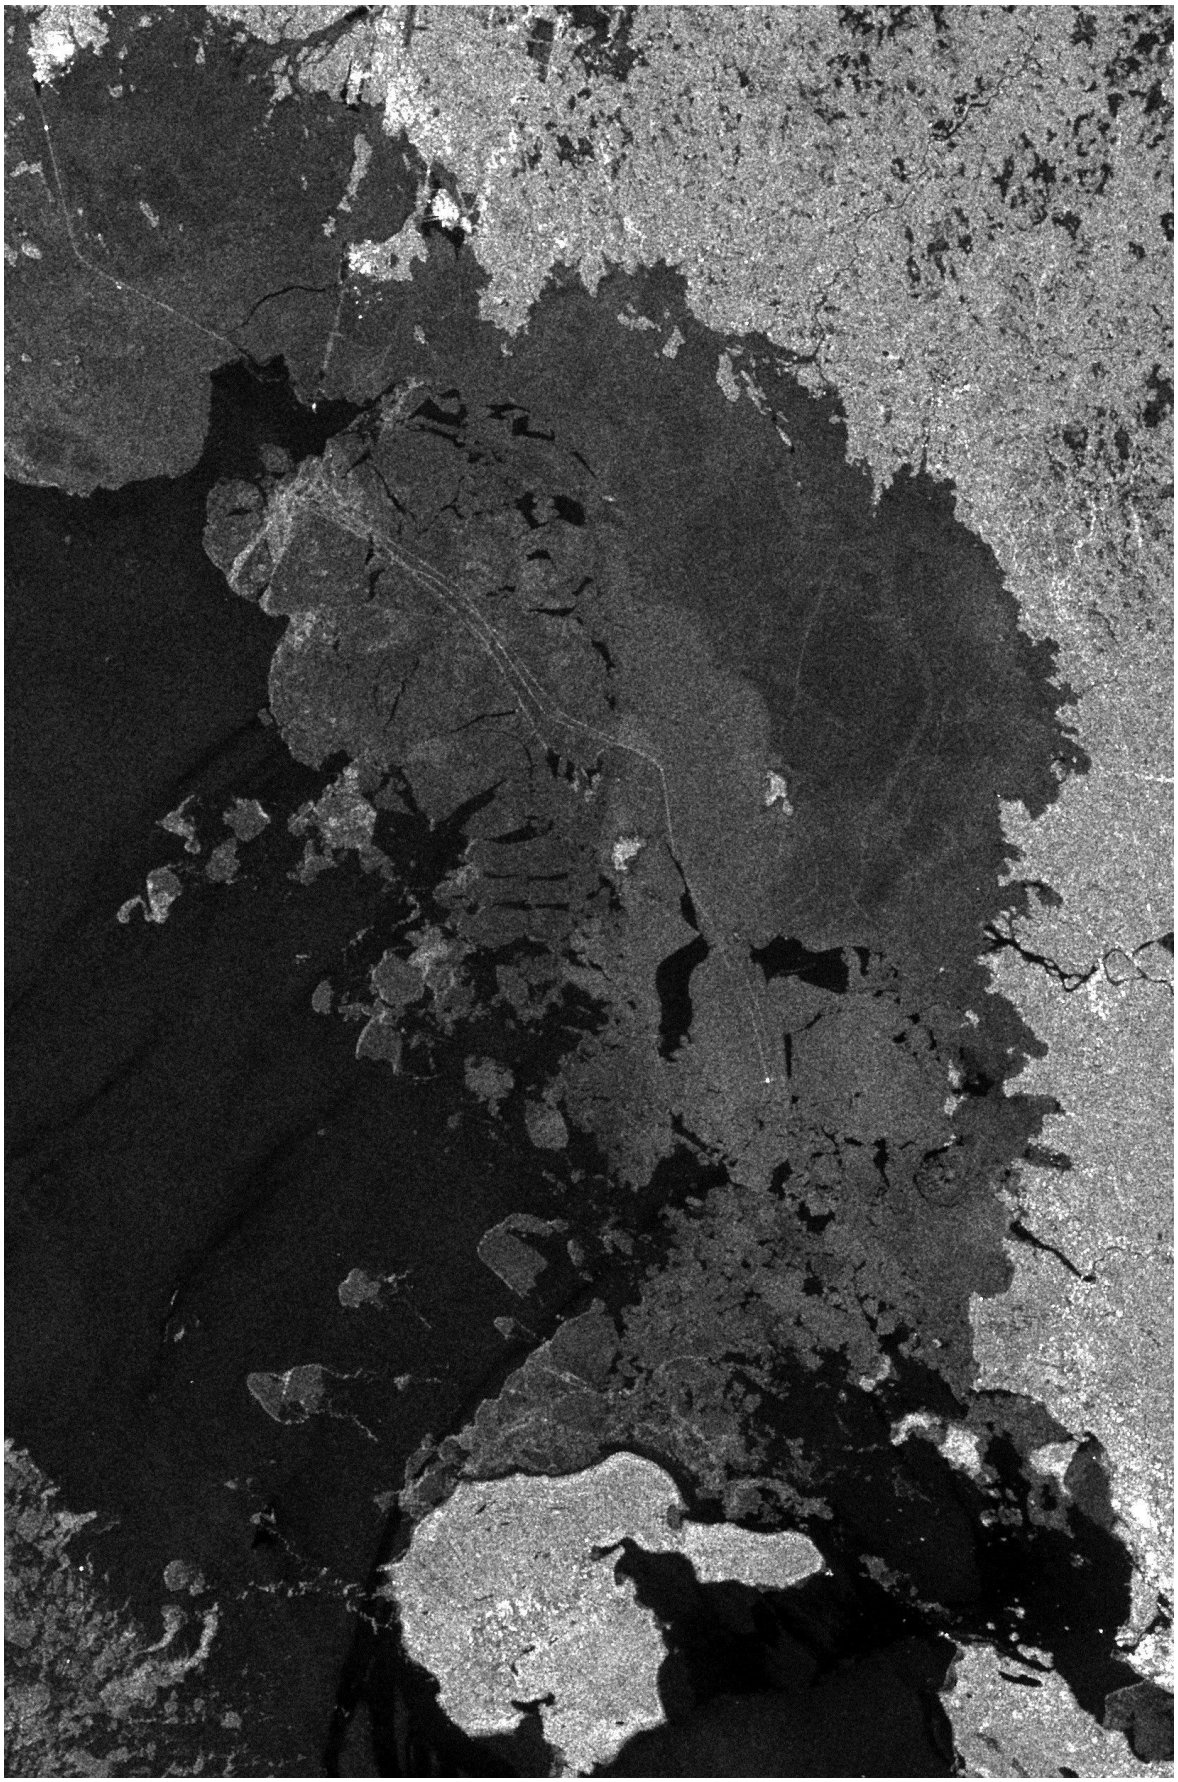

Figure 3-24. Bay of Bothnia, 21 Apr 2015. Last image from the series. For the first time, the track parallel to the coast is split into several parallel ones.
S1A_EW_GRDM_1SDH_20150421T154936_20150421T155000_005584_007250_6531

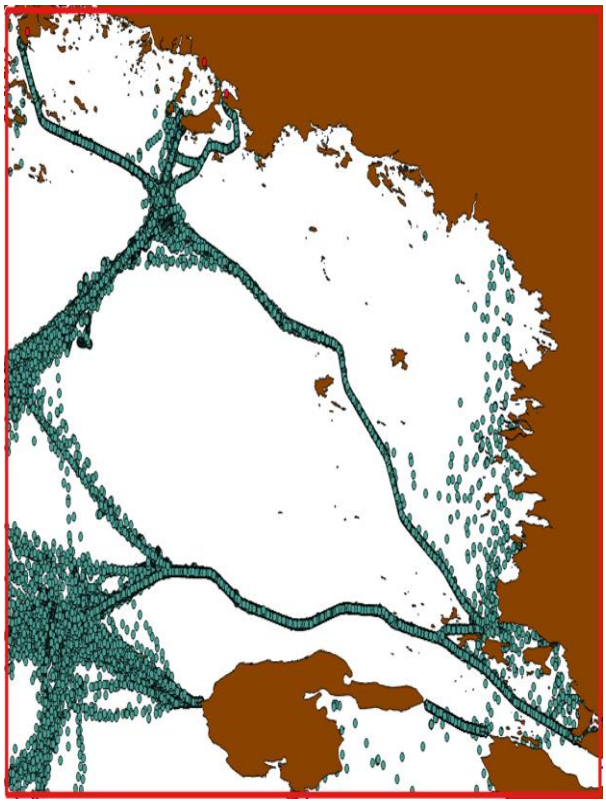

Figure 3-25. AIS message positions from 1 Jan to 22 Apr 2015; 3.6 months, ~645,000 messages. The AIS positions coincide with the bright tracks in the Sentinel-1 images. Except on the open water in the bottom left, and in an area to the right in front of the coast. The MMSI number of the latter AIS positions shows the ship to be a hovercraft, which can move over the ice instead of through it.

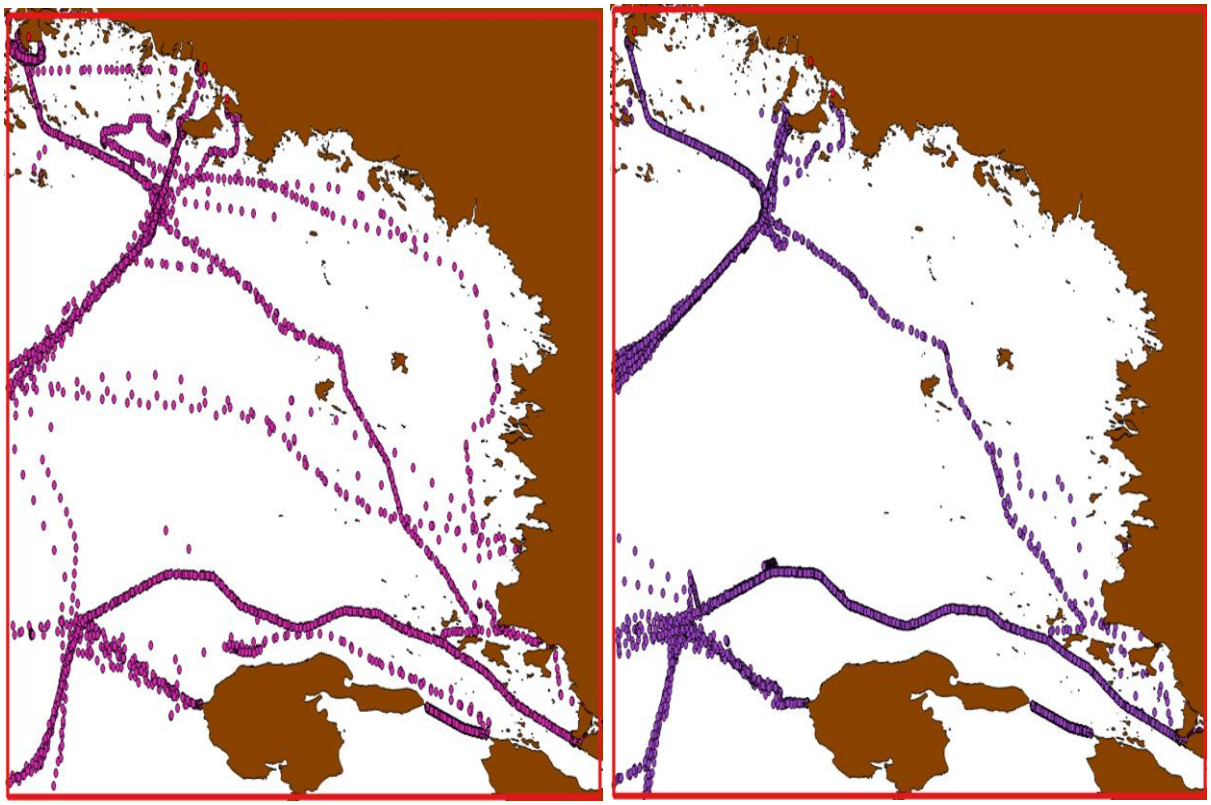

Figure 3-26. AIS message positions for one week. Left, 6-12 Oct 2014, before freezing (49k messages). Right, 2-8 Feb 2015, while sea ice was present (47k messages).

4 Sentinel-1 in the Western Indian Ocean

A total of 1,874 Sentinel-1 images of the Western Indian Ocean region have been analysed with JRC's ship detection software (SUMO). The images were acquired between 3 October 2014 and 15 November 2015, i.e. over a period of 1.1 year. All the Sentinel-1 images released by ESA for that area during that period were used in the campaign. A breakdown of the images by mode and product is given in Table 4-1, and their location footprint in Figure 4-1. All the products were Ground Range Detected (GRD) and included single- and dual-polarisation images. The goals of this campaign were:

1. To support the PMAR-MASE project, which aimed to provide experience to the authorities in the region on maritime surveillance. See [G2015] and [PMAR2015] for more details on this project. During the course of the project, ship reporting data (AIS – terrestrial and satellite – and LRIT) were used to continuously update the Maritime Situational Picture in the whole region. Satellite SAR images where acquired and analysed to get information on the non-reporting ships in the area.
2. To prepare JRC Maritime Affairs Unit's IT infrastructure and processes to deal with the large amounts of Sentinel-1 images that have started to arrive. All of the images were downloaded from ESA's Sentinel Hub [SHub2015] using scripts and were processed in fully-automatic batch mode with SUMO. Some of the images were additionally processed in a semi-automatic mode, with human refining of the results.
3. To learn about Sentinel-1's capabilities for maritime surveillance.

For the PMAR-MASE project, also images of a few other satellite SAR sensors have been analysed. They are not discussed here, but some results are included in [G2015]. The PMAR-MASE project had a somewhat smaller area of interest than what was used for the Sentinel-1 analysis, being between 31 and 68 degrees longitude and between -30 and 19 degrees latitude (outlined in Figure 4-1).

Table 4-1. Overview of Sentinel-1 images analysed in the Western Indian Ocean campaign.

Mode	Product	# Images	Swath width (km)	Resolution (m)
EW	GRDH	14	400	50
	GRDM	339	400	90
IW	GRDH	1,467	250	20
SM	GRDH	54	80	23
Total		1,874		

The ship detections from SUMO were correlated with the positions of known ships obtained from interpolating received ship position reports to the times of the satellite images. Ship position reports were received from:

- Satellite AIS:
  - Over the whole region of Figure 4-1 up to 5 Nov 2015: from the Norwegian Coastal Administration (2-3 AIS receiving satellites).
  - Inside the PMAR-MASE box up to 15 Sep 2015: from exactEarth, Orbcomm/LuxSpace and SpaceQuest (up to 14 AIS receiving satellites). After 15 Sep 2015 (end of the PMAR-MASE satellite AIS contract), the incoming AIS data started to taper off, leaving only 3 satellites of SpaceQuest until the end of the Sentinel-1 campaign.
- Coastal AIS:
  - From the MSSIS network: over the whole region and throughout the entire campaign.
  - From exactEarth and Orbcomm/LuxSpace: as for satellite AIS.
- LRIT from ships that use the EU LRIT Data Centre ("EU-Flagged ships"): inside the PMAR-MASE box and up to 15 Sep 2015.

Therefore, during the last two months of the campaign but especially in the last 10 days, there was significantly less data from known ships available.

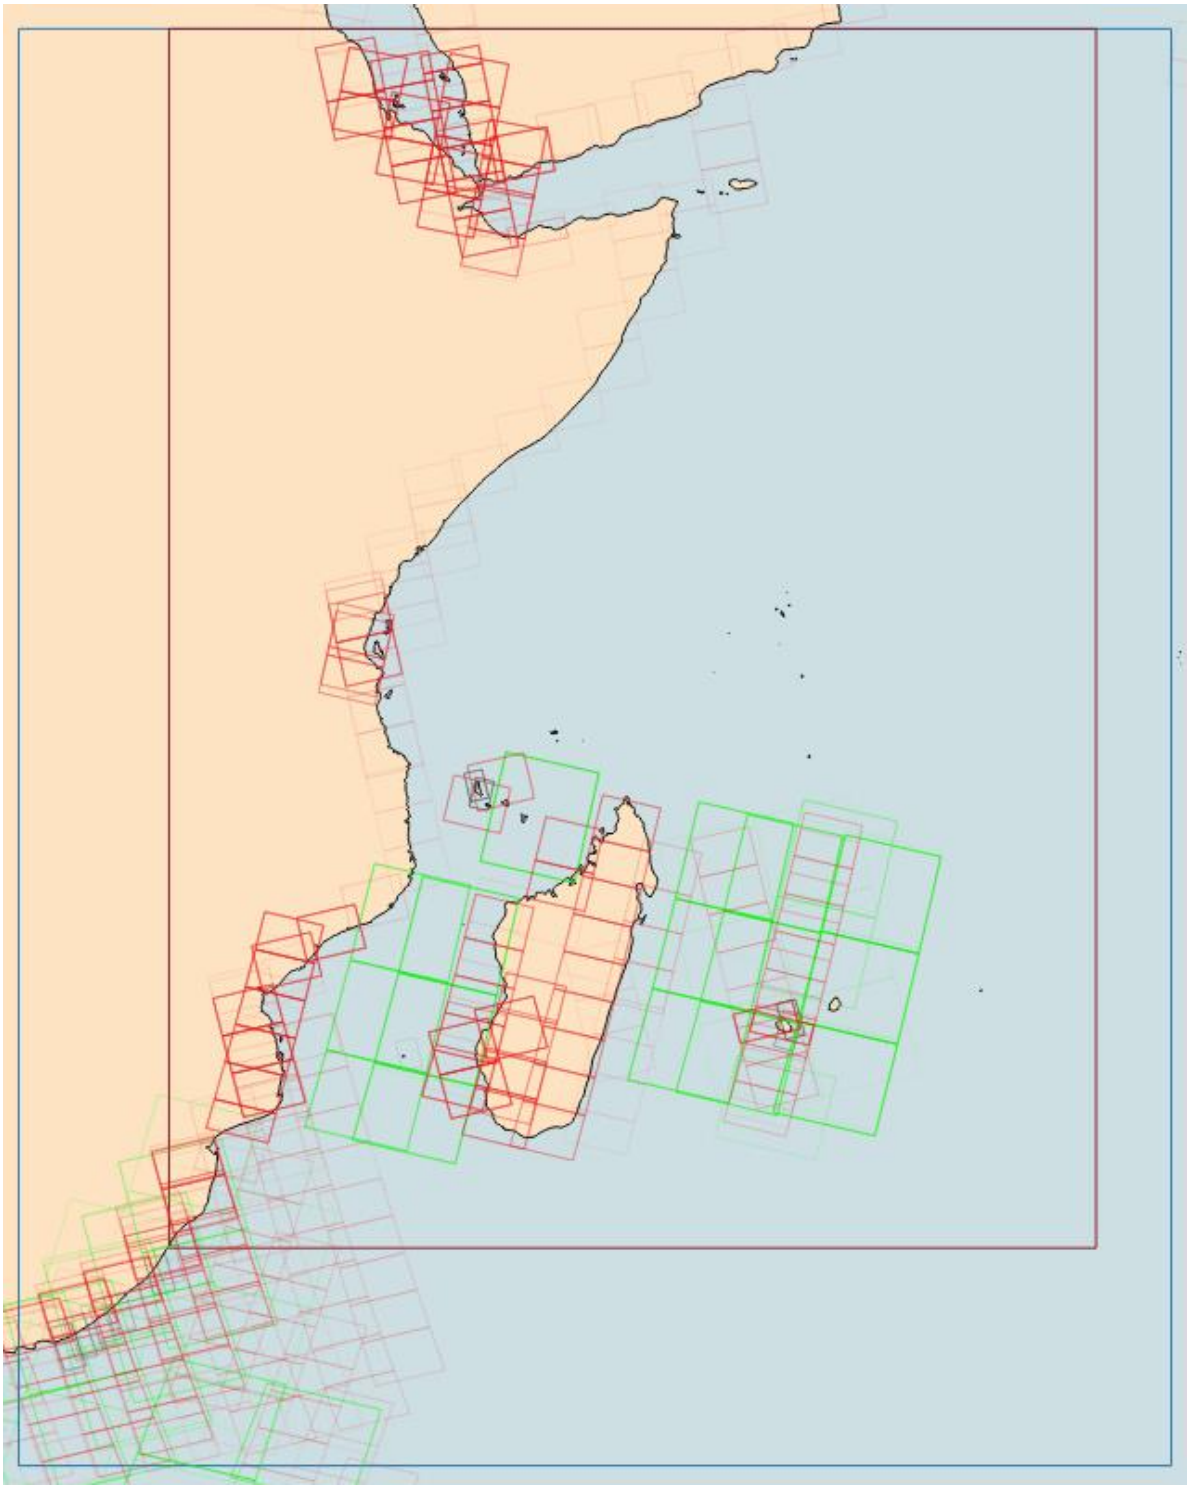

Figure 4-1. Western Indian Ocean campaign area (external blue box) and footprint of all 1,874 Sentinel-1 images (black: SM mode; red: IW mode; green: EW mode). A darker frame indicates multiple images with the same footprint. PMAR-MASE area indicated with brown box.

4.1 Ship detection in batch mode

This section presents the results of the fully-automatic processing of the 1,874 Sentinel-1 images. Since the images were going to be analysed without the possibility for a human to verify the individual detection results, a number of SUMO parameters were adjusted so that the results didn't include too many false alarms. These adjustments were:

- The detection thresholds were raised. For some frequently occurring sea clutter types, most false alarms are small, weak targets, so by excluding the weaker targets, most false alarms are avoided. The two most problematic sea clutter types are (a) waves that occur at high wind speeds, and (b) the edges of areas of extremely low wind, and therefore extremely low radar cross section (RCS). This also meant that some weak targets were not detected. The threshold value in the co-polarised channels (HH and VV) needed to be raised by a large amount to prevent many false alarms from these types of sea clutter. The threshold in the cross-polarised channels (HV and VH) was only marginally raised, and then mainly to prevent the detection of false alarms in the border between sub-swaths EW1 and EW2, and between EW4 and EW5. Examples of very low wind areas and sub-swath borders that often result in false alarms are given in Figure 4-2.
- OpenStreetMap data (layer = coastline), buffered by 250 metres, was used to mask out the land in the ship detection analysis. The mask was manually expanded to exclude from the analysis coastal areas in Madagascar, Kenya, Tanzania and the port of Maputo that contain reefs or wetland that are unlikely to include ships but that typically give rise to many false alarms if included in the analysis.

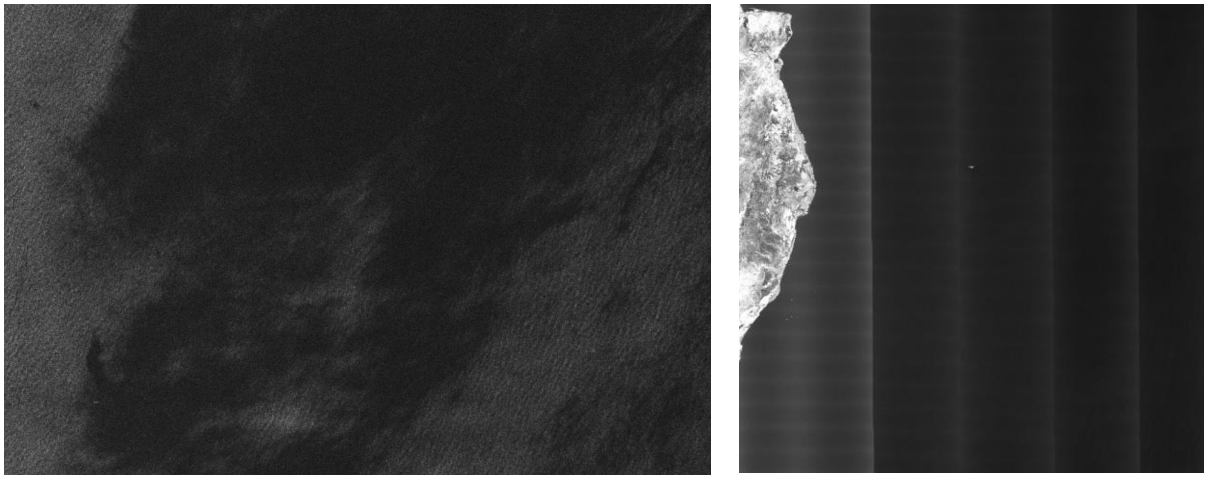

Figure 4-2. Patterns in Sentinel-1 images that tend to lead to false alarms if standard detection thresholds are employed. Left: area of low RCS in a VV-pol IW image. False alarms tend to occur at the edges of the dark areas. Right: interfaces between sub-swaths in a VH-pol EW image, especially abrupt between EW1 and EW2, and EW4 and EW5 (EW1 is on the left, EW5 on the right).
S1A_IW_GRDH_1SDV_20151113T164440_20151113T164505_008588_00C2EF_A507
S1A_EW_GRDM_1SDV_20151103T024257_20151103T024401_008434_00BEB7_6EB5

The correlation between the SUMO ship detections and the reporting ships was also done automatically, on the Blue Hub.

The total number of detections (VDS targets) was 24,121, although 2,601 of those were (automatically) classified by SUMO as azimuth ambiguities or other azimuth artefacts. Of the 21,520 non-ambiguity detections, 12,409 (that is 58 %) were correlated to reporting ships. Therefore, 42 % of the detections in the SAR images could not be correlated to reporting ships. Table 4-2 shows the breakdown by image mode and product. It is not possible to know exactly how many of the detections correspond to real ships and how many to false alarm, but using their recurrence as an indication of their likelihood of being a false alarm [SG2015], it can be estimated that at least 2,679 (735 correlated, 1,944 uncorrelated) of the detections are false alarms, most of them originating

from reefs or wetland outside Madagascar, Kenya, Tanzania and the port of Maputo, and a few of them being unrecognized azimuth or range ambiguities. Subtracting the recurrent targets from the number of detections and from the number of correlated detections, it would lead to a total of 18,841 detections of actual ships, 11,674 of which are correlated to reporting ships. With this new estimate, it can be concluded that about 38 % of the SAR detections likely to be actual ships could not be correlated to reporting ships. Figure 4-3 presents a classification tree with these numbers normalised. The way to read the diagram is as follows: for every 100 SAR detections (at the top) there are 11 azimuth ambiguities, 11 recurrent targets (likely to be false alarms), 48 detections likely to be actual ships and correlated to reporting ships, and 30 detections likely to be actual ships and not correlated to reporting ships. Also, 38 % of the detections likely to be actual ships are not correlated to reporting ships.

Starting from the side of the reporting ships, of the total number of 19,597 interpolated reporting ship positions (that is, the positions of the reporting ships that were within the SAR images' extents at the image acquisition times), 3,244 were inside the buffered land mask used by SUMO (most of them are stationary ships in ports) and therefore would never be detected. Of the remaining 16,353 interpolated points, 11,674 (that is 71 %) were correlated to non-recurrent SAR detections. This percentage (71 %) can be interpreted as an estimate of the detection rate of reporting ships (outside the buffered land mask) achieved by SUMO in the Sentinel-1 images during this campaign. Since non-reporting ships tend to be smaller (and therefore more difficult to detect) than reporting ones, the achievable detection rate of non-reporting ships is expected to be below 71 %. The diagram in Figure 4-4 is to be interpreted as follows: for every 100 interpolated reporting ship positions, 17 are inside the buffered land mask and cannot be detected in SAR images, 20 are outside the buffer and uncorrelated, 4 are correlated to recurrent SAR detections, and 59 are correlated to non-recurrent SAR detections. Also, of the interpolated points outside the buffer, 71 % are correlated to non-recurrent detections. In other words, 29 % of the reporting ships outside the buffered land mask have not been detected in the SAR images.

It has to be remembered that these numbers are just estimates, and that several sources of uncertainty exist: false SAR detections that are not ambiguities or recurrent, land masking inaccuracies, incorrectly interpolated points, incorrectly correlated targets, etc.

Figure 4-5, Figure 4-6 and Figure 4-7 present the geographical distribution of all 21,520 SAR detections, the 12,409 detections that have been correlated to reporting ships, and the 9,111 detections that have not been correlated, respectively.

Table 4-2. Ship detection and correlation results per mode and product in the Sentinel-1 Western Indian Ocean campaign.

Mode	Product type	# Images	# VDS targets	# Interpolations	# Correlations	# VDS targets uncorrelated	% VDS targets uncorrelated	# Interpolations uncorrelated	% Interpolations uncorrelated
EW	GRDH	14	552	546	406	146	26	140	26
	GRDM	339	6,456	5,006	3,456	3,000	46	1,550	31
IW	GRDH	1,467	14,037	13,830	8,403	5,634	40	5,427	39
SM	GRDH	54	475	215	144	331	70	71	33
Total		1,874	21,520	19,597	12,409	9,111	42	7,188	37

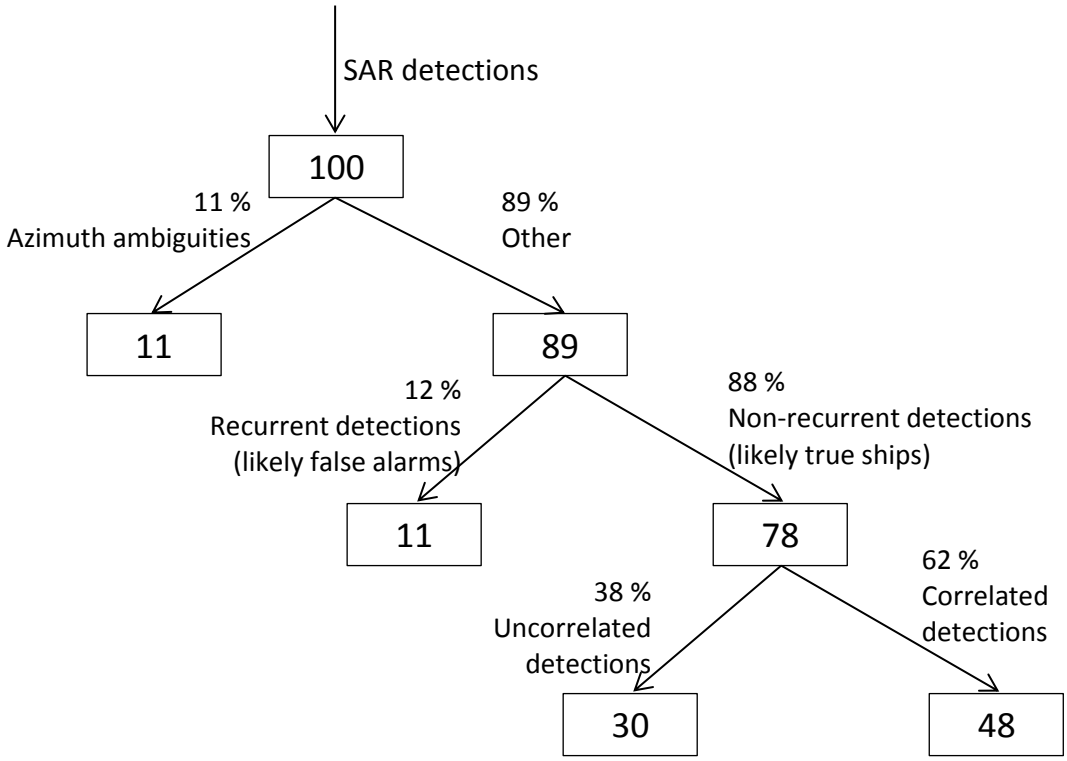

Figure 4-3. Classification tree of SAR detections in the Western Indian Ocean campaign. Figures in the boxes indicate number of detections in that particular branch of the tree, starting from a normalised value of 100 at the top.

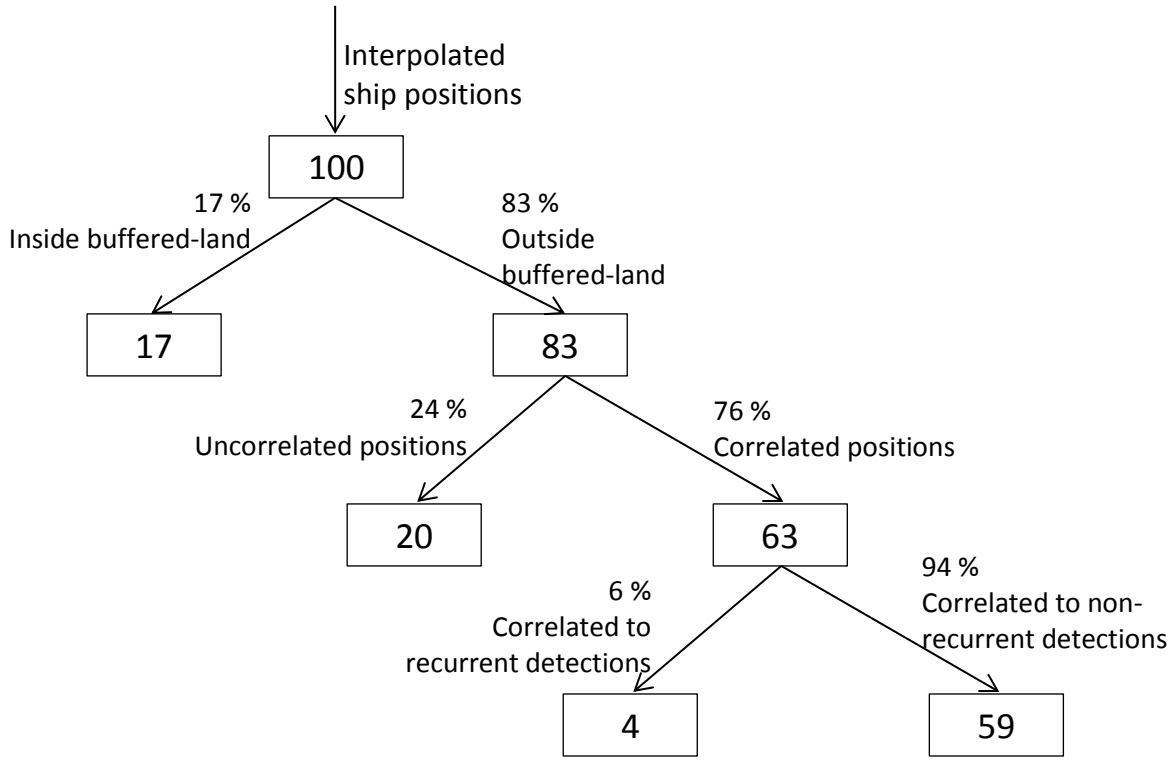

Figure 4-4. Classification tree of interpolated reporting ship positions in the Western Indian Ocean campaign. Figures in the boxes indicate number of detections in that particular branch of the tree, starting from a normalised value of 100 at the top.

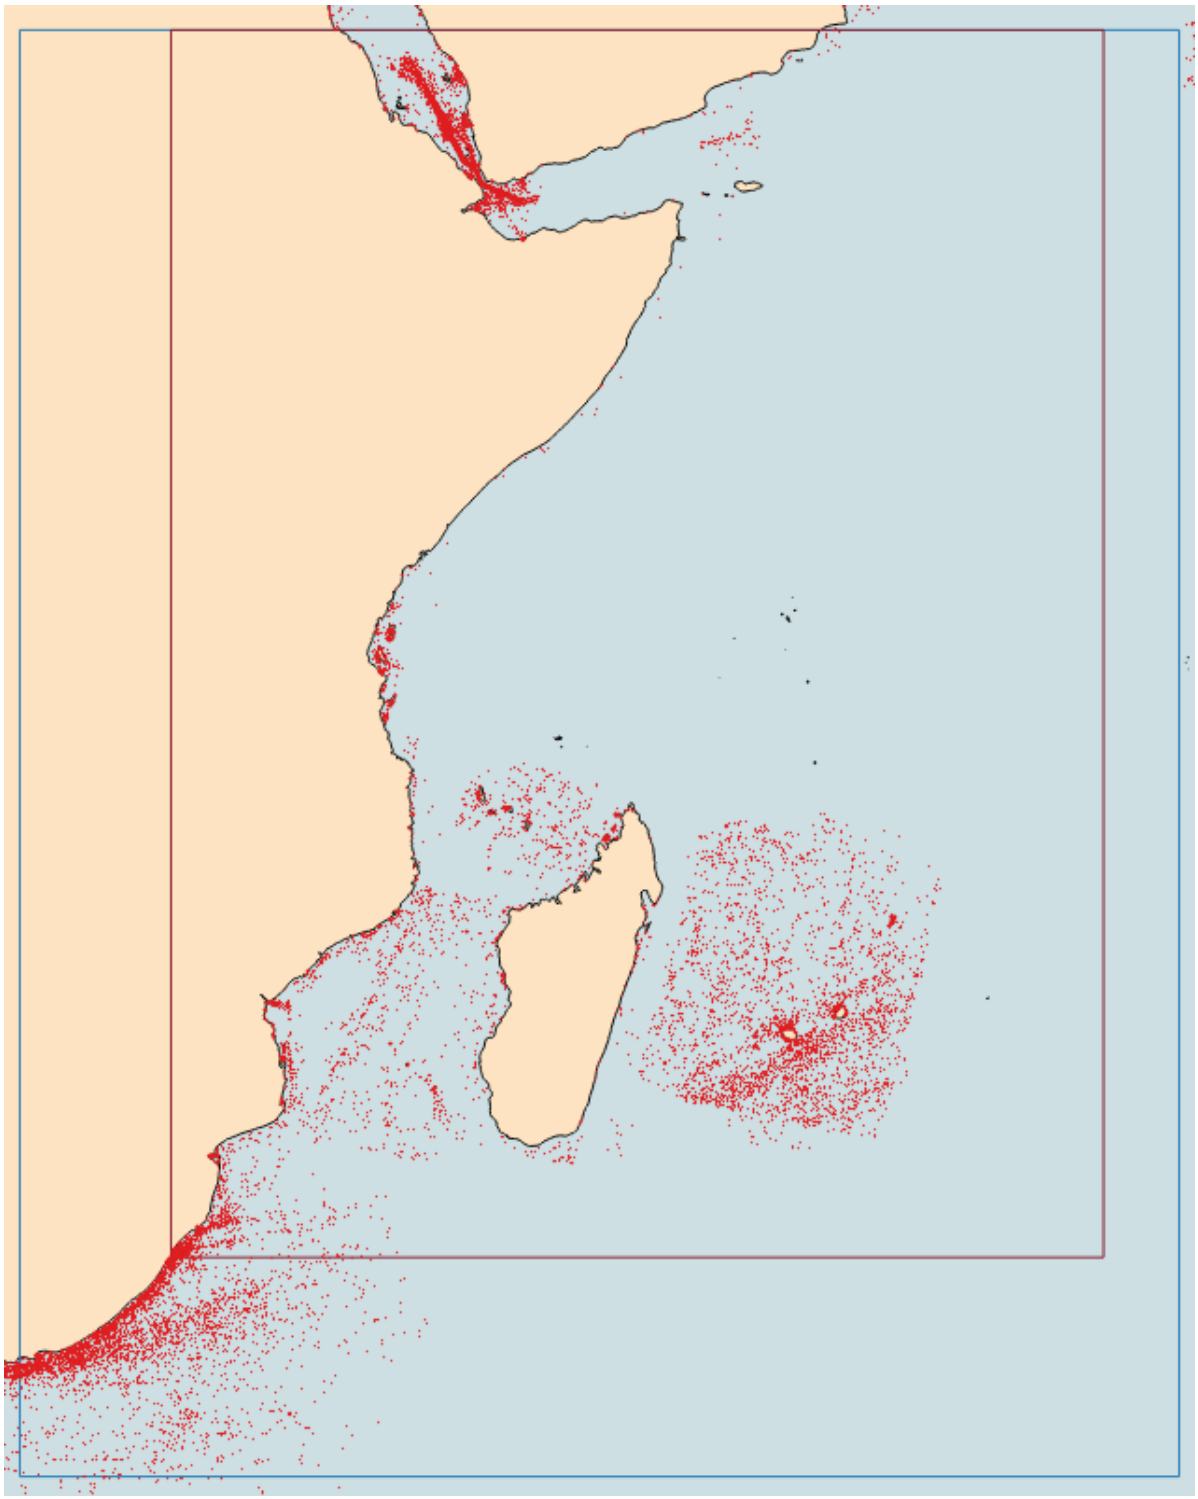

Figure 4-5. Geographical distribution of all 21,520 Sentinel-1 SAR detections in the Western Indian Ocean campaign.

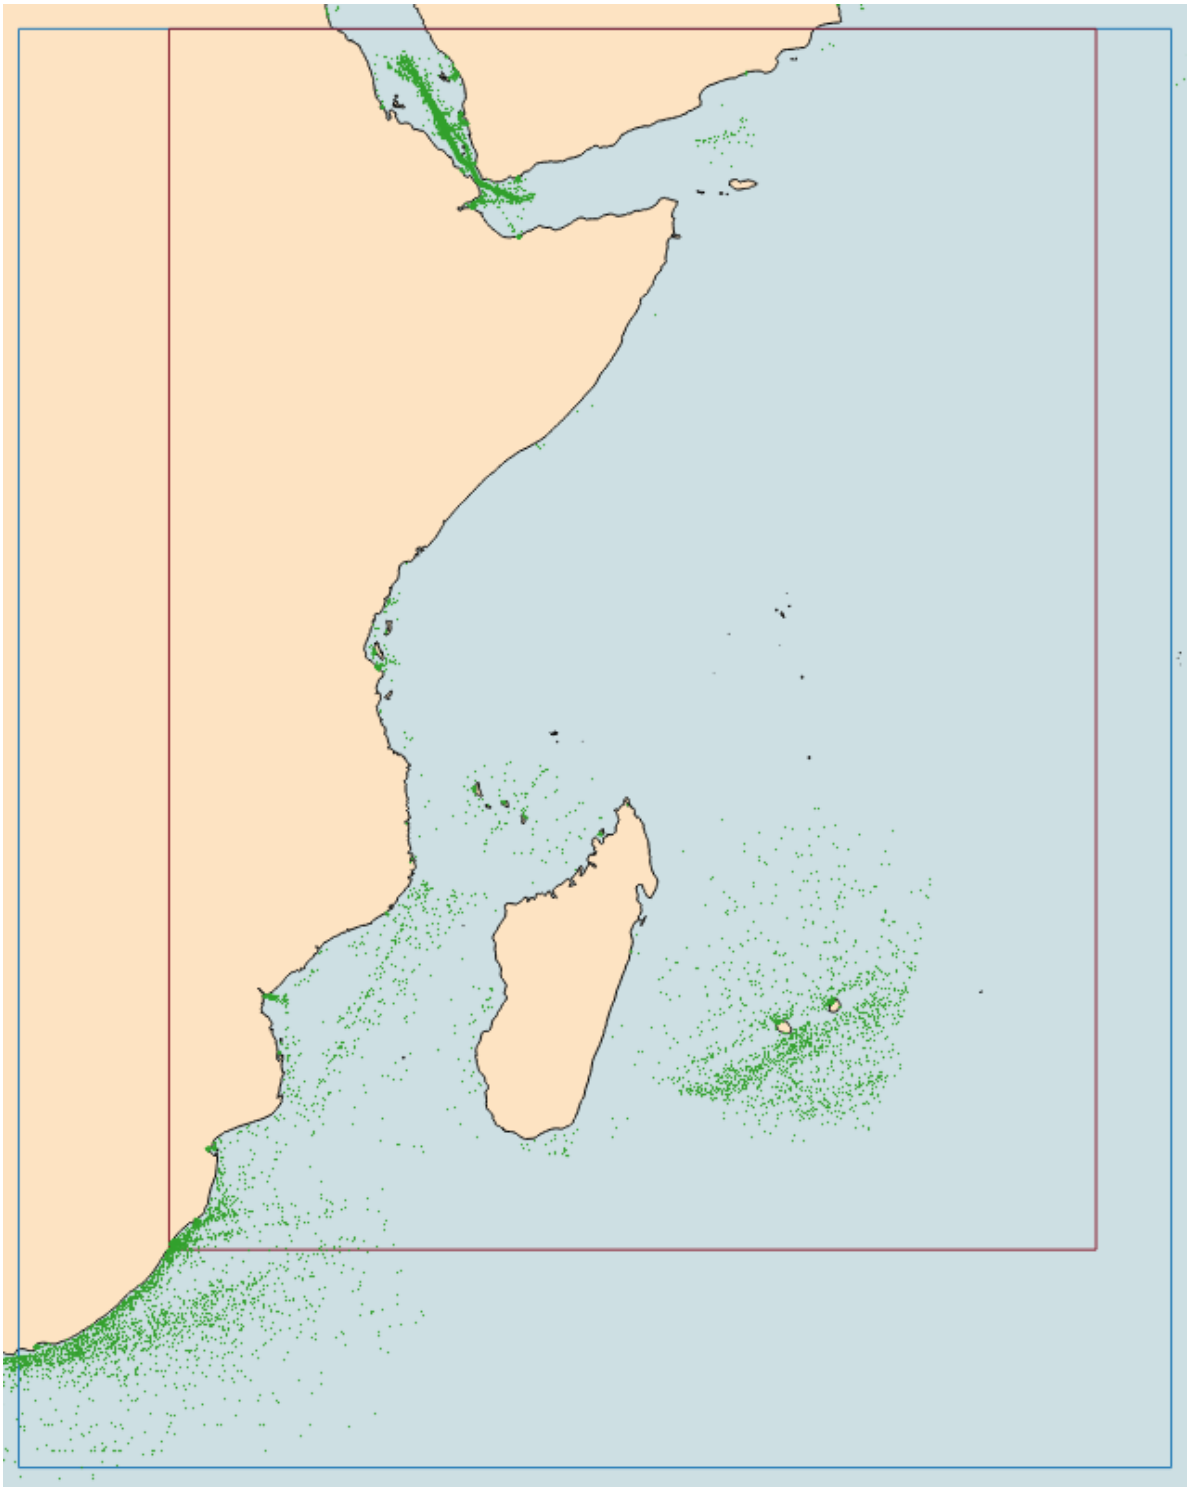

Figure 4-6. Geographical distribution of the 12,409 Sentinel-1 SAR detections that could be correlated to reporting ships.

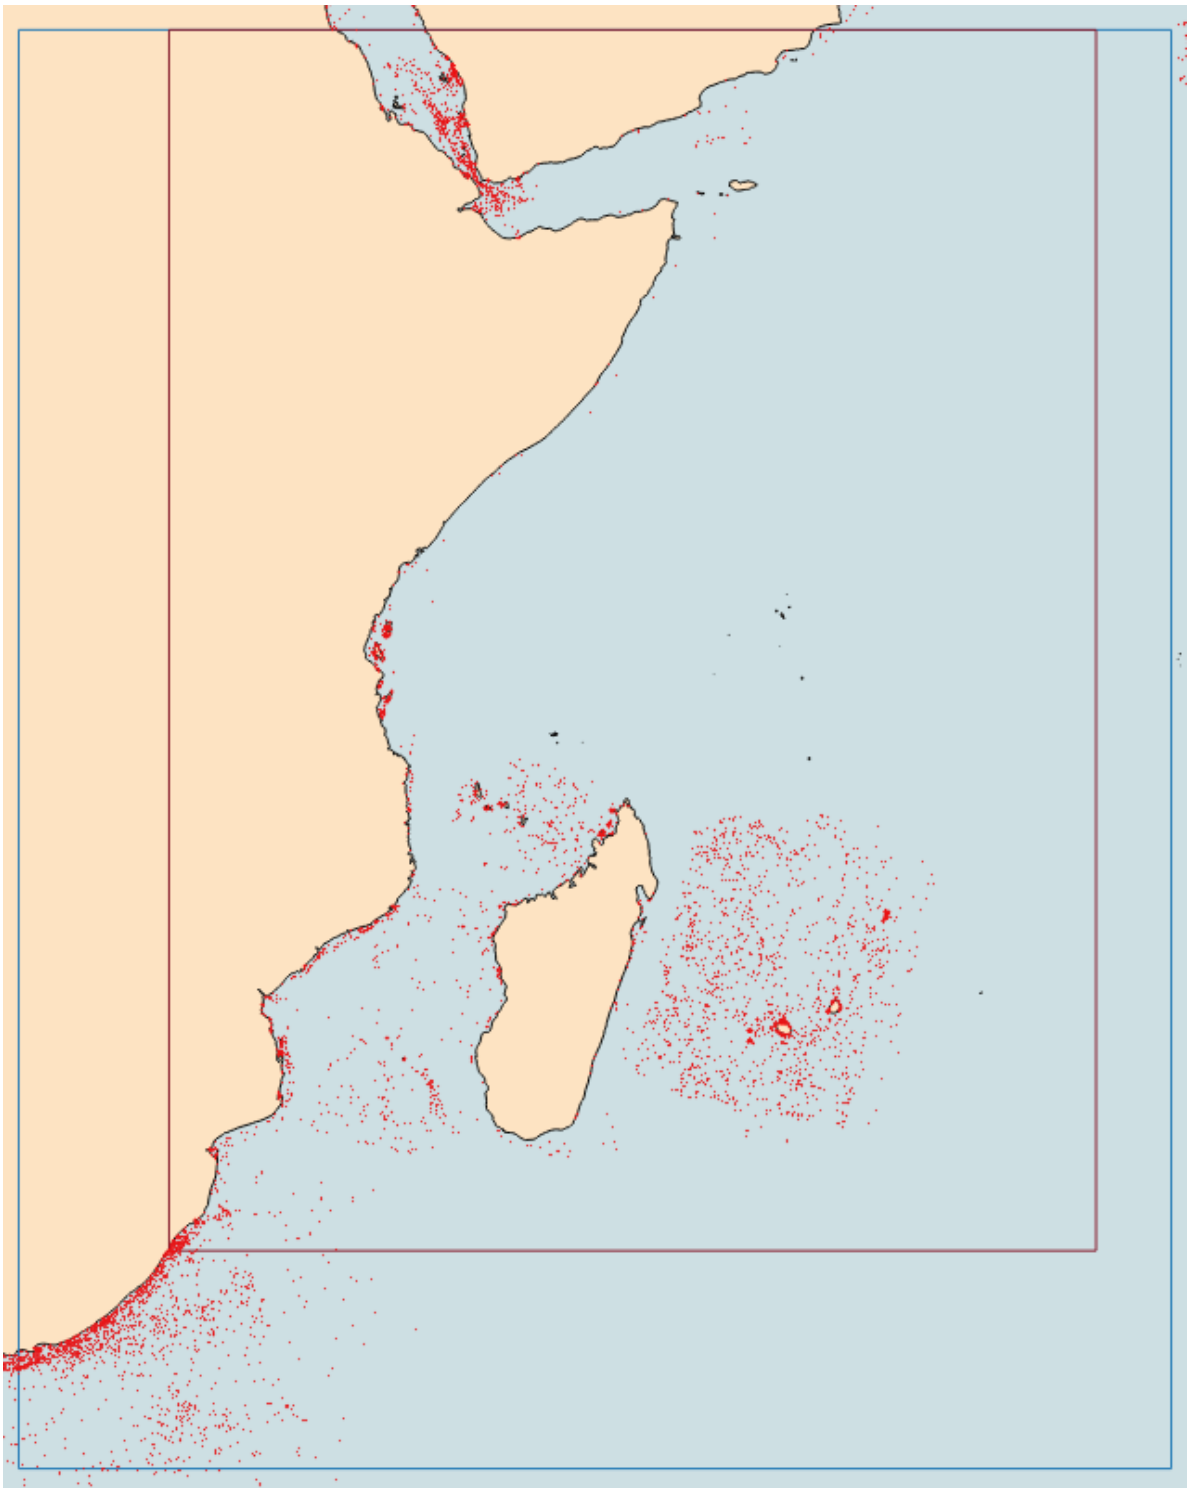

Figure 4-7. Geographical distribution of the 9,111 Sentinel-1 SAR detections that could not be correlated to reporting ships.

5 Sentinel-1 in the Mediterranean Sea

A total of 1,557 Sentinel-1 images of the Mediterranean Sea have also been analysed with JRC's ship detection software (SUMO). The images were acquired from 3 October 2014 to 11 September 2015, with most of the images being from 2014 or beginning of 2015. 1,394 of the images were of the IW mode, 140 were EW (35 GRDH and 105 GRDM), and 23 were SM. All the products were Ground Range Detected (GRD) and included single- and dual-polarisation images. Ship reporting data, used for correlation between Sentinel-1 detections and known ships, were mainly restricted to terrestrial AIS (MSSIS) and some satellite AIS. During the campaigns, terrestrial AIS coverage was good along the coasts of Spain, northern Italy (both the Ligurian and Adriatic seas), western Greece, Turkey, Cyprus and Malta, and generally poor in the rest of the areas. An estimate of the MSSIS coverage is shown in Figure 5-1. It is therefore acknowledged that the completeness of reporting data is poorer in this campaign than in the Western Indian Ocean campaign.

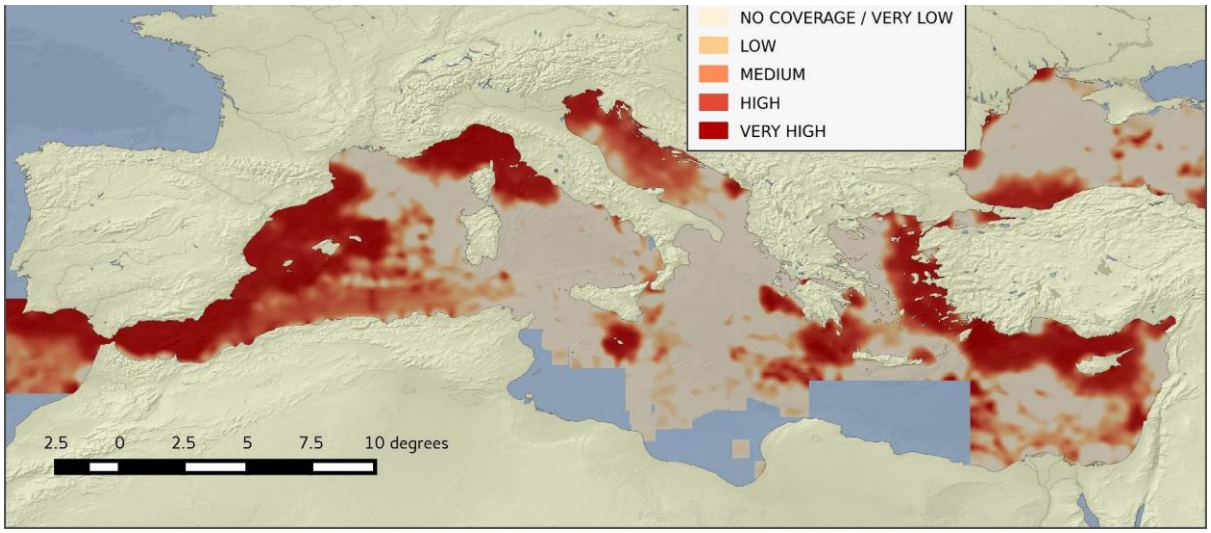

Figure 5-1. Estimated coverage of the MSSIS coastal AIS network for the month of January 2015. This was computed on the Blue Hub [AA2015].

The overall ship detection and correlation results are reported in Table 5-1. From this table, it is clear that the Med is much more densely navigated than the Western Indian Ocean (92,019 against 21,520 in terms of SAR detections, for a similar number of images, and 102,982 against 19,597 in terms of interpolated reporting ship positions). The percentage of SAR detections that have not been correlated to reporting ships is 49 %, compared to values of 38 % - 42 % in the Indian Ocean. The higher figure in the Med was expected as a result of the poorer AIS coverage in this area. The percentage of interpolated positions left uncorrelated is 54 %, compared to the 29 % - 37 % estimates in the Indian Ocean. The higher value in the Med can probably be explained by the fact that in this sea there is a higher proportion of small vessels (fishing ships, sailing boats, etc.) that transmit AIS than in the Indian Ocean. These smaller vessels are harder to detect in SAR images.

Figure 5-2 shows the Sentinel-1 images analysed in this campaign, the 92,019 SAR detections, the 47,373 detections that have been correlated to reporting ships, and the 44,646 detections that have not been correlated. The high number of detections allows distinguishing the main traffic routes in the Med. The correlations map also indicates where the AIS coverage was good. The high density of detections off the coast of Israel and positioned in lines near-parallel to the east-west direction are the result of blocky image artefacts that lead to false alarms in the ship detection process. An example of these blocky artefacts is shown in Figure 5-3. The cause of the artefacts could be related to signal interference from the ground.

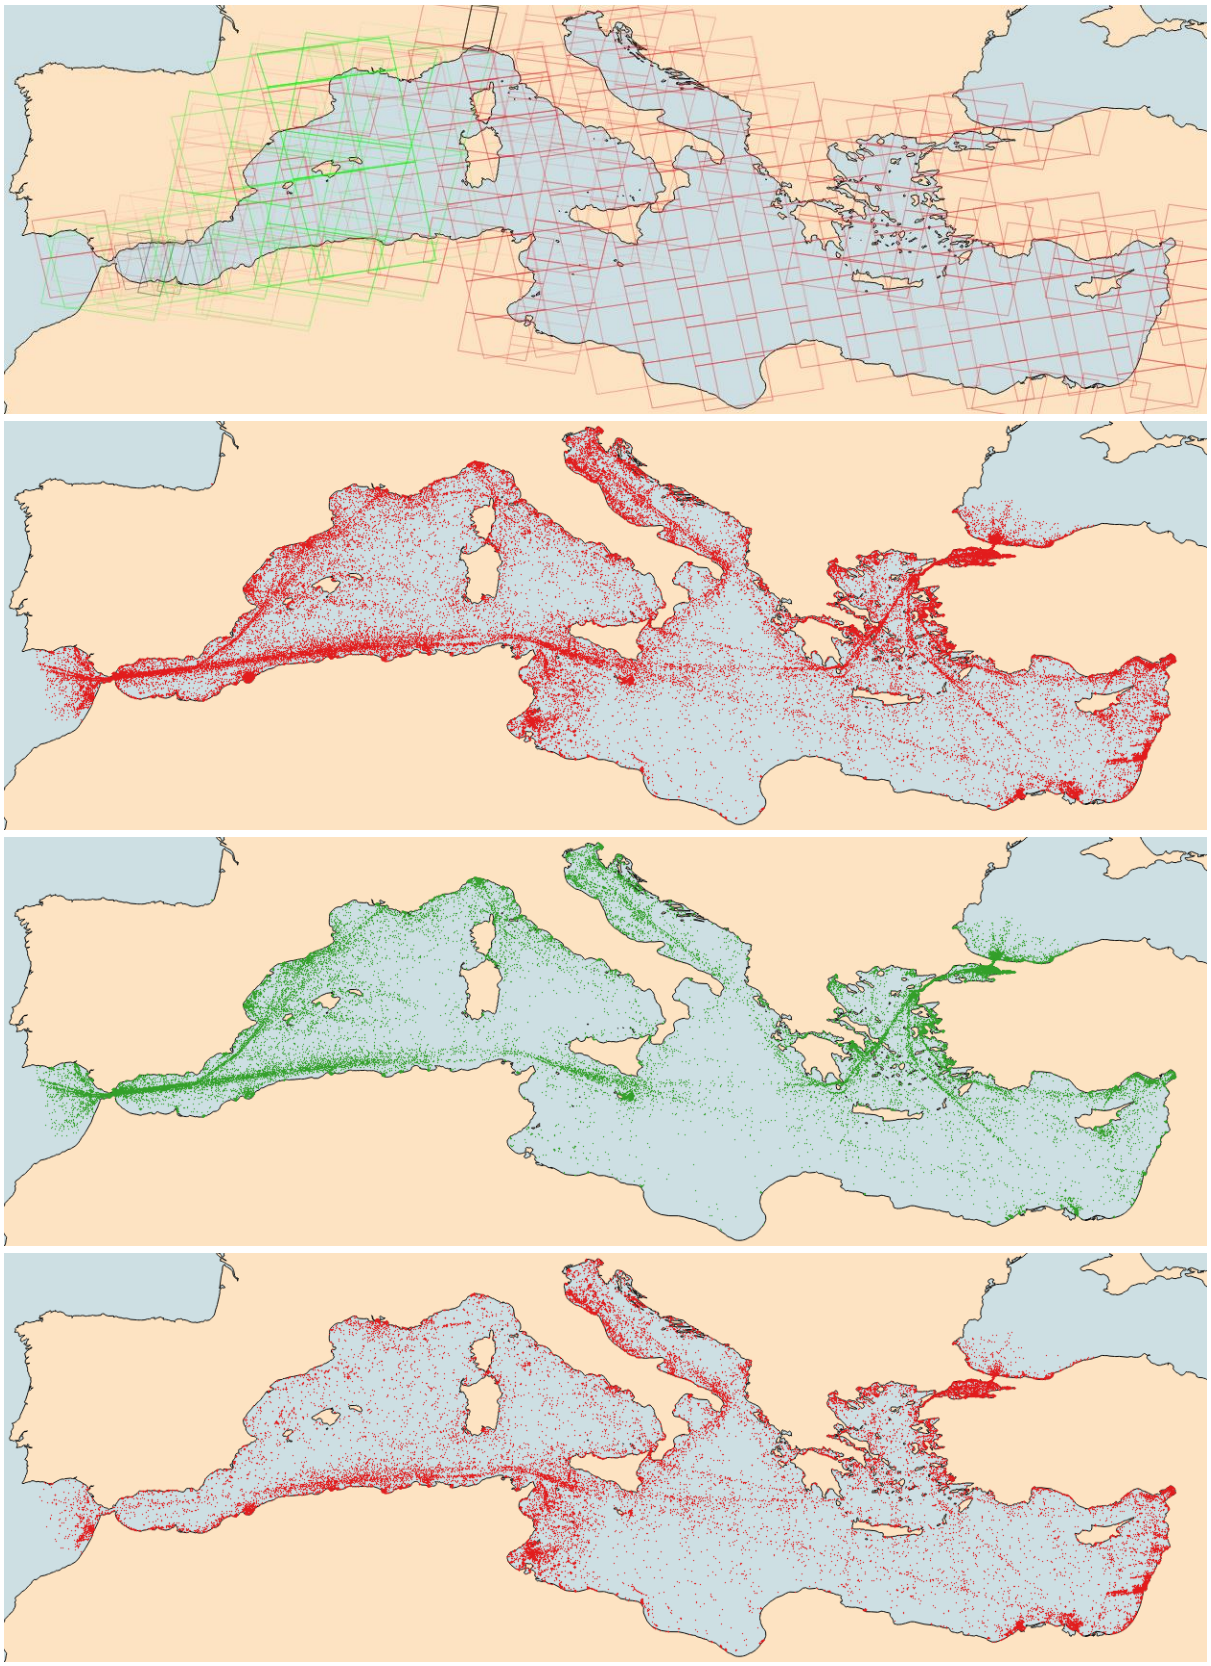

Figure 5-2. From top to bottom: Sentinel-1 images (1,557) analysed in this campaign (SM in black, IW in red, EW in green, a darker frame indicates multiple images with the same footprint); all SAR detections (92,019); detections that have been correlated to reporting ships (47,373); detections that have not been correlated to reporting ships (44,646).

Table 5-1. Ship detection and correlation results per mode and product in the Sentinel-1 Med Sea campaign.

	#	%	%
# Images	1,557		
# SAR detections (VDS targets) including azimuth artefacts	98,187		
# SAR detections (excluding azimuth artefacts)	92,019	100	
# Interpolated reporting ship positions	102,982		100
# Correlations	47,373	51	46
# VDS targets uncorrelated	44,646	49	
# Interpolations uncorrelated	55,609		54

The recurrent target analysis has been carried out for this data set. The set-up for this analysis is as follows: a recurrent target is deemed to exist if at least two SAR detections occur in the same geographical position (to within 50 meters) in different images. Only the geographical position is taken into account to determine recurrence; other detection features (length, strength…) are not assessed. A total of 15,828 SAR detections have been found to be recurrent (4,109 recurrent targets). Following visual examination of these recurrent targets, they can be classified as:

1. Azimuth or range ambiguities (or other types of image artefacts) that had not already been removed in SUMO's azimuth ambiguity rejection step. Many more of this type of recurrent target would have been detected had the usual ship detection threshold for the co-polarisation images been used.
2. Recurrent detections outside ports. Often they are caused by mooring points, or the ships moored to them.
3. Actual fixed offshore structures (most likely linked to oil exploration, off the coast of Tunisia, Libya, Egypt, in the Adriatic Sea, etc.) and support ships around the structures that hold the same position for a period of time.

Unlike in the Western Indian Ocean campaign, there do not appear to be many recurrent targets resulting from land masking accuracies (reefs, wetlands, etc.)

Some of the recurrent targets have been correlated to reporting ships. This is normally the case of ships stationed outside ports (if AIS coverage is available in that location) and, sometimes, the support vessels around oil platforms. Other recurrent targets have not been correlated. This is generally the case for the ambiguities. Figure 5-4 presents the geographical distribution of the recurrent SAR detections that have been correlated and not correlated, and the non-recurrent detections, correlated and not correlated.

Figure 5-3. Blocky artefacts off the coast of Israel that result in false ship detections. S1A_IW_GRDH_1SDV_20141023T154024_20141023T154049_002959_0035CE_F4DD

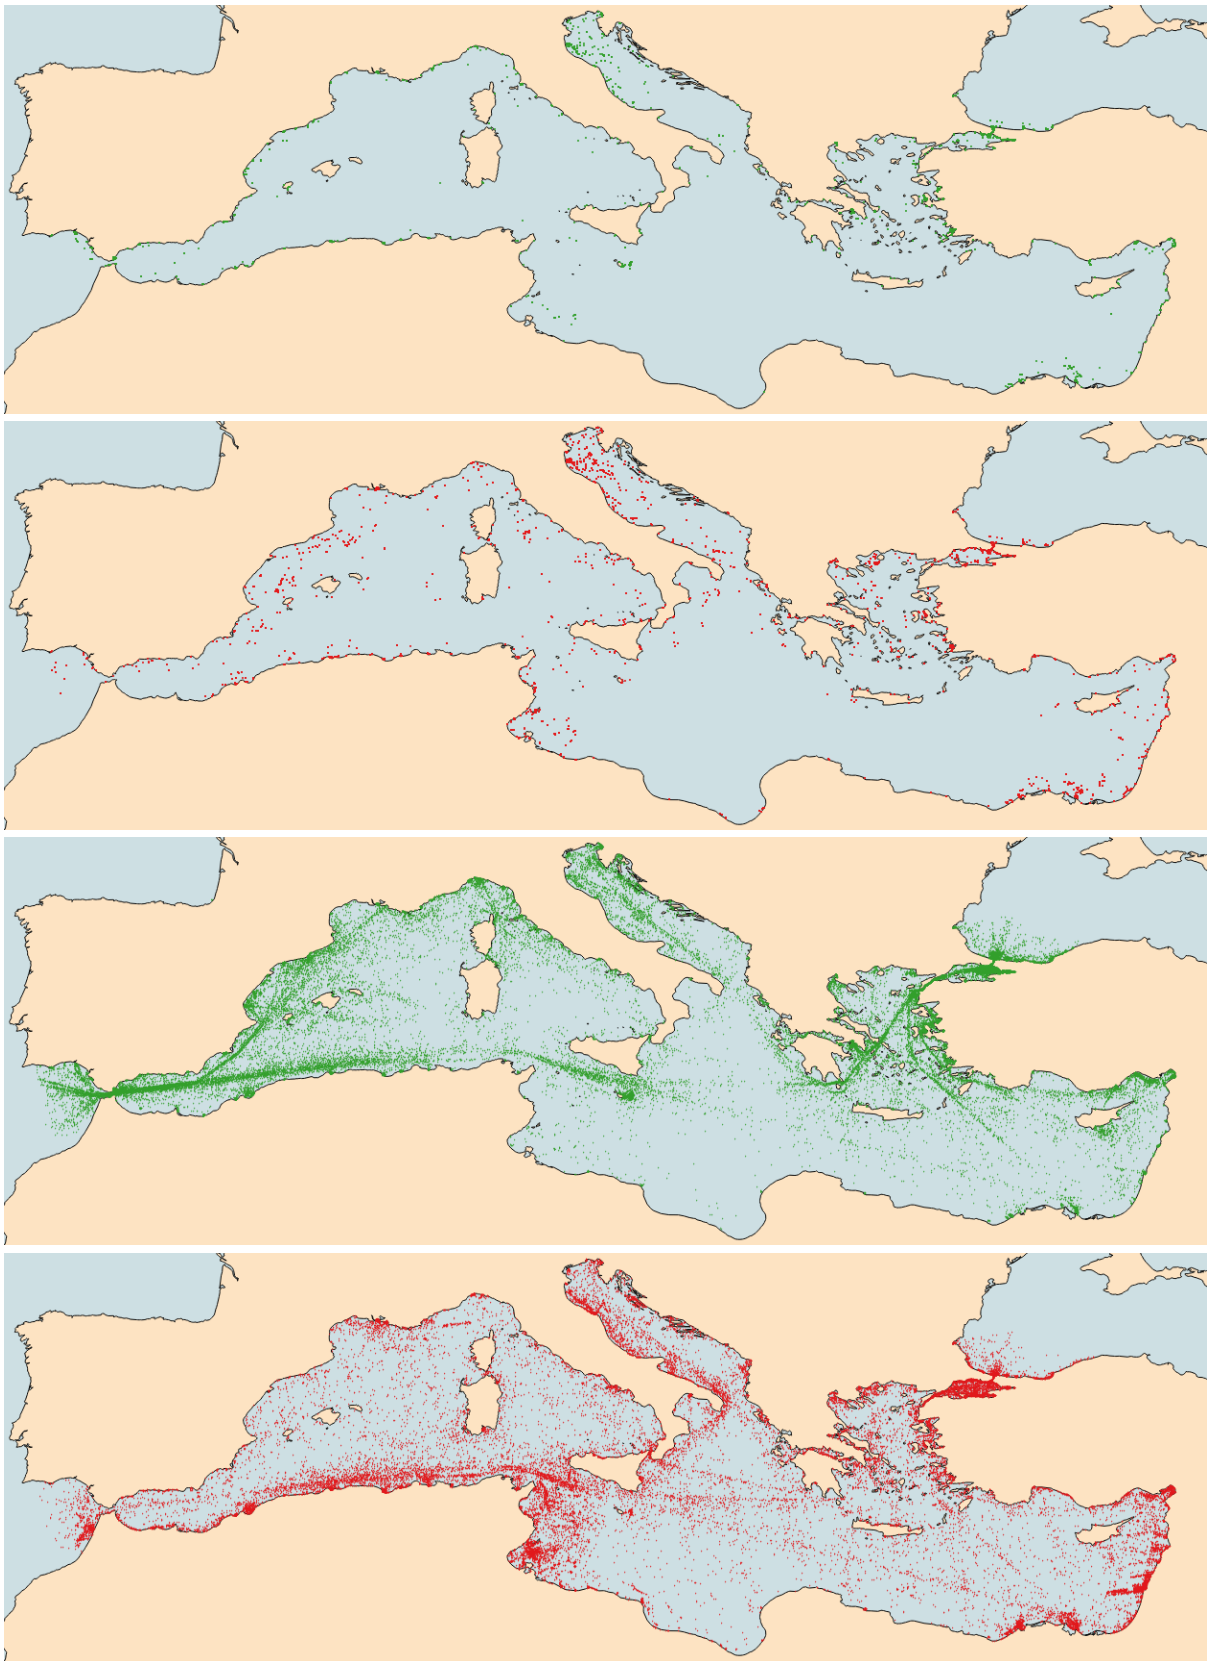

Figure 5-4. From top to bottom: recurrent SAR detections correlated to reporting ships (5,445); recurrent SAR detections not correlated (10,383); non-recurrent SAR detections correlated (41,928); non-recurrent SAR detections not correlated (34,263).

6 Ship size estimation in Sentinel-1 images

This chapter focuses on the analysis of Sentinel-1 images for the automatic estimation of vessels' length and width. Four different estimation methods have been applied to selected images from the Western Indian Ocean campaign (Chapter 4) and then compared to provide a detailed performance assessment.

6.1 Dataset

In order to create a reliable ground truth (GT) of vessels for the benchmarking, a subset of 4 images have been selected from the Western Indian Ocean campaign. As the same dataset had to be used also for the vessel classification presented in Chapter 7, all the images have been chosen among those acquired in the IW mode, which provides a spatial resolution of about 20 m, the only one available to try to distinguish the ship types.

Hereafter, the overview images and details of the images are reported. (The quoted dimensions are in pixels, and the incidence angles are at mid-range.)

- *S1A_IW_GRDH_1SDH_20141108T162905_20141108T162930_003192_003AD3_87BF*

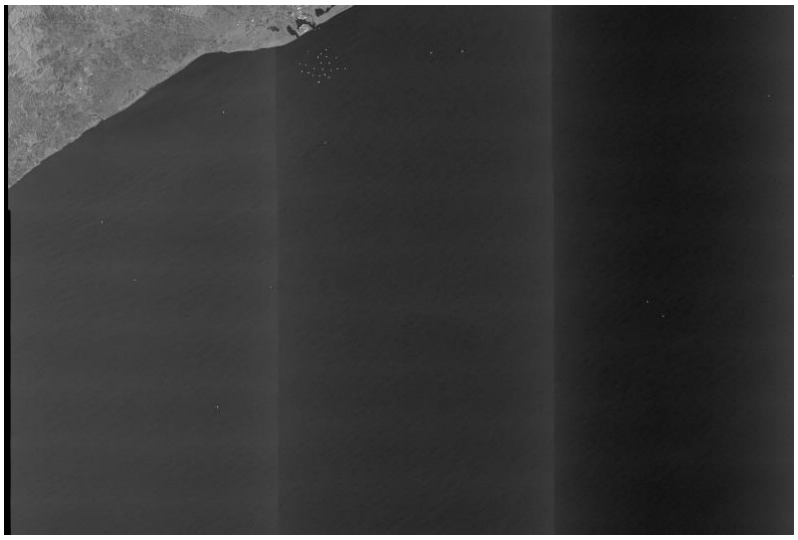

Image	20141108T162905
Sensor	Sentinel-1
Product	GRDH
Mode	IW
Polarization	HH/HV
Start Time	2014-11-08 16:29:05
Stop Time	2014-11-08 16:29:30
Pass	Ascending
Looking Direction	Right
Dimensions	25606 x 16884
Rg x Az Res [m]	20 x 22
Inc. Angle [deg]	38.5

- *S1A_IW_GRDH_1SSV_20150622T151820_20150622T151845_006488_0089A8_9BD7*

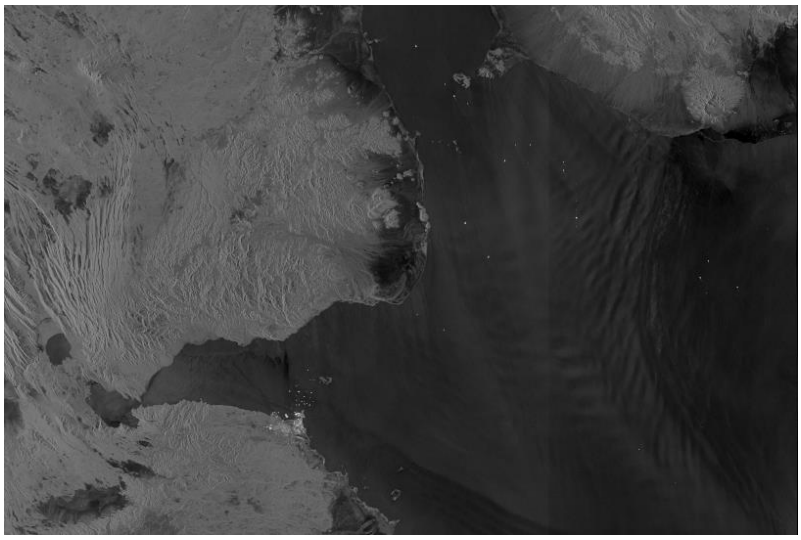

Image	20150622T151820
Sensor	Sentinel-1
Product	GRDH
Mode	IW
Polarization	VV
Start Time	2015-06-22 15:18:20
Stop Time	2015-06-22 15:18:45
Pass	Ascending
Looking Direction	Right
Dimensions	25149 x 16810
Rg x Az Res [m]	20 x 22
Inc. Angle [deg]	39.1

- *S1A_IW_GRDH_1SSV_20150821T151840_20150821T151908_007363_00A1DE_612C*

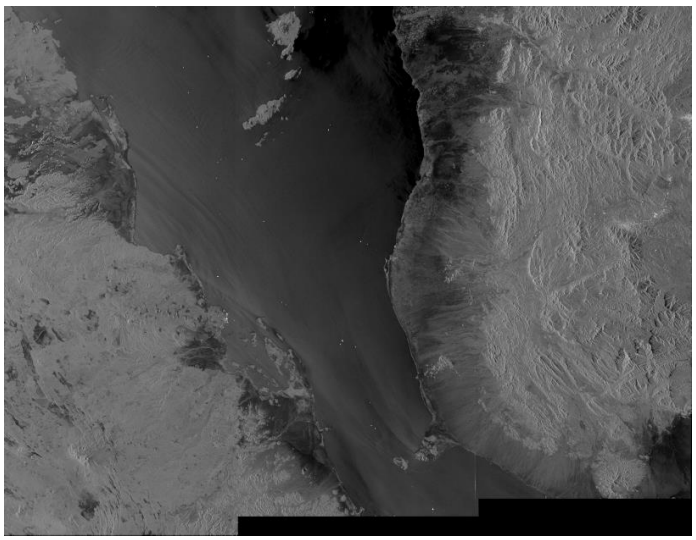

Image	20150821T151840
Sensor	Sentinel-1
Product	GRDH
Mode	IW
Polarization	VV
Start Time	2015-08-21 15:18:40
Stop Time	2015-08-21 15:19:08
Pass	Ascending
Looking Direction	Right
Dimensions	25144 x 19414
Rg x Az Res [m]	20 x 22
Inc. Angle [deg]	39.0

- *S1A_IW_GRDH_1SSV_20150821T151908_20150821T151933_007363_00A1DE_9F42*

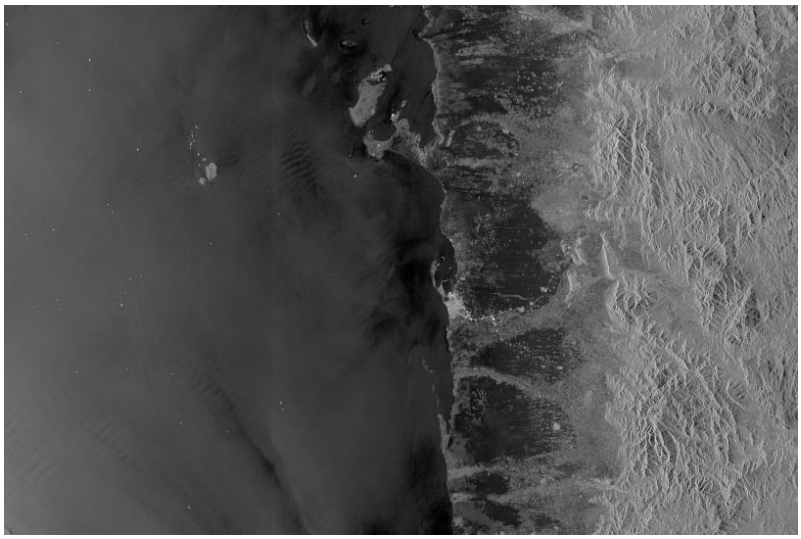

Image	20150821T151908
Sensor	Sentinel-1
Product	GRDH
Mode	IW
Polarization	VV
Start Time	2015-08-21 15:19:08
Stop Time	2015-08-21 15:19:33
Pass	Ascending
Looking Direction	Right
Dimensions	25145 x 16792
Rg x Az Res [m]	20 x 22
Inc. Angle [deg]	39.0

The VDS targets detected in batch mode (i.e., in an unsupervised way) by SUMO have been then correlated on the Blue Hub to the AIS data available from that area and visually checked to ensure the accuracy of the GT.

Details about this correlation and validation process are reported in Table 6-1. The total number of VDS target detected within the 4 images was 146: 120 of them were correlated to AIS reporting vessels. Out of 146, the targets that have been visually confirmed to be both well correlated and with reliable AIS information (in some cases, the relevant parts of the AIS message were missing so it was not possible to use them) were 107.

Table 6-1. VDS detection and correlation to AIS (#T = number of targets; #C = number of correlated targets; #VT = number of validated targets).

Image	Campaign	UTC	Run version	#T	#C	#VT
20141108T162905	Western Indian Ocean	2014-11-08T16:29:05	BATCH_MODE	38	33	29
20150622T151820	Western Indian Ocean	2015-06-22T15:18:20	BATCH_MODE	46	36	31
20150821T15184	Western Indian Ocean	2015-08-21T15:18:40	BATCH_MODE	26	24	23
20150821T151908	Western Indian Ocean	2015-08-21T15:19:08	BATCH_MODE	36	27	24
				146	**120**	**107**

6.2 Results

For the benchmarking, each image of the dataset has been analysed with four different estimation algorithms: a MATLAB code developed in the past years and that has been widely used to analyse previous campaigns (MTO); the one currently implemented in SUMO (MTS); a new algorithm based on mathematical morphology developed last year (MTN) [S2014]; and a further refinement for width reconstruction that has been introduced to improve the results of MTN (MTNN).

Table 6-2 and the histogram in Figure 6-1 show the mean relative error in length and width estimation for the whole dataset. As regards the length estimation, the 4 methods are almost equivalent, as the average error for all of them is around 20 %. The best results are achieved by the new improved algorithm MTNN, whose error is 16.8 %, about 3 percentage points (pp) less than MTS.

The major improvements introduced by MTNN are immediately clear looking at the bars that show the width reconstruction error. The results for MTNN are significantly better not only with respect to its previous version MTN (100 pp), but also with respect the current SUMO estimation algorithm (more than 100 pp) and MTO (more than 10 pp). In absolute terms (see Figure 6-2), these errors correspond to an average mismatch in length that is similar for the 4 algorithms (comprised between 32 m and 40 m), whereas, as far as width is concerned, the error ranges from 10 m (MTNN) to 50 m (MTS).

To better understand these numbers, Figure 6-3 and Figure 6-4 report the scatterplot of the estimated lengths and widths versus the actual ones. The black lines correspond to the ideal behaviour.

As regards the lengths, the scatterplot confirms that MTNN is the best method (MTN, as expected, is quite similar), as the dispersion of the values is less significant than the others. Moreover, its behaviour seems not to be biased by underestimation or overestimation, as it happens for MTO and MTS. We can in fact notice that the blue line is generally below the black theoretical line (meaning that the MTO algorithm is underestimating lengths, especially for bigger ships), whereas the slope and intercepts of the purple line show in the case of MTS there is an overestimation of smaller ships and an underestimation of bigger ones.

Concerning the widths, we can see that, In general, all the MTO, MTN and MTS are not behaving well. They actually have slopes similar to the theoretical one, but they are all above the black line, so they are all overestimating widths. Differently, MTNN has a less steep slope and it tends to overestimate smaller ships and to underestimate larger ships. Nevertheless, the dispersion of the green points is quite limited and the estimations provided are more in line with the actual widths.

As a general comment, it is worth noting that most of the ships of the dataset have breadths around 30 m (the smallest is 12 m wide, the largest 60 m wide), thus, at the given resolution, targets have widths of a few pixels. This intrinsically leads to rough estimations of the widths that result in large relative errors.

Table 6-2. Mean relative and absolute errors (entire dataset).

Average Error	MTO	MTN	MTNN	MTS
Length (%)	19.2	18.3	16.8	19.4
Width (%)	55.2	132.6	33.6	174.2
Length [m]	39.8	33.8	31.9	35.9
Width [m]	16.5	37.3	10.4	50.6

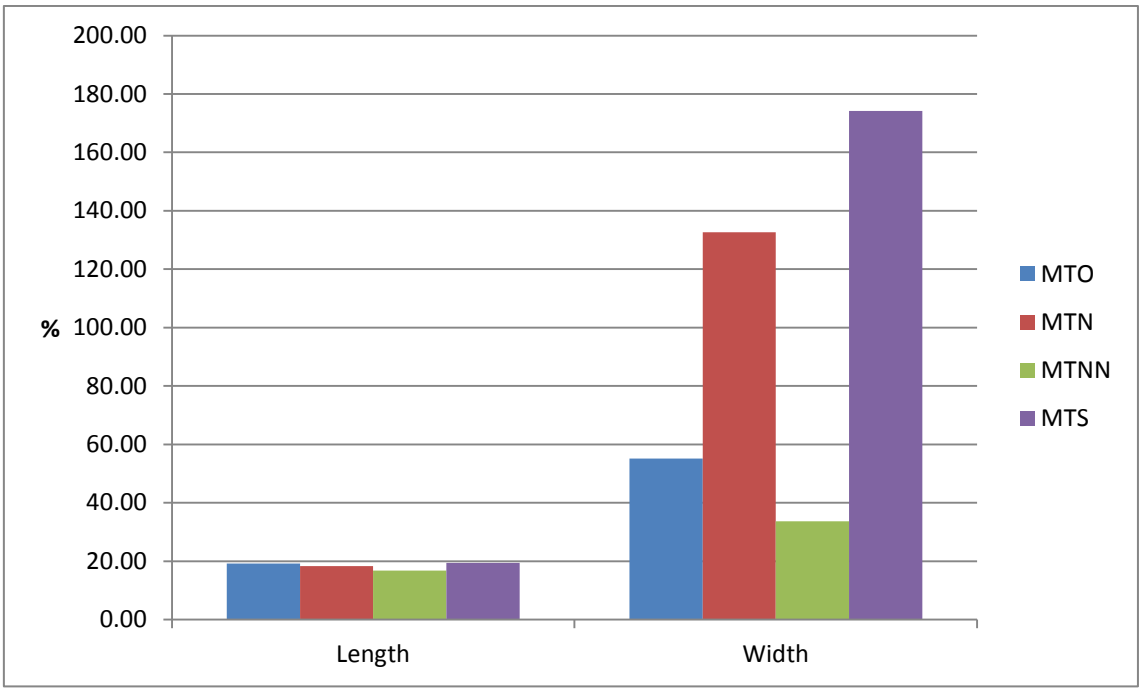

Figure 6-1. Mean relative error (entire dataset).

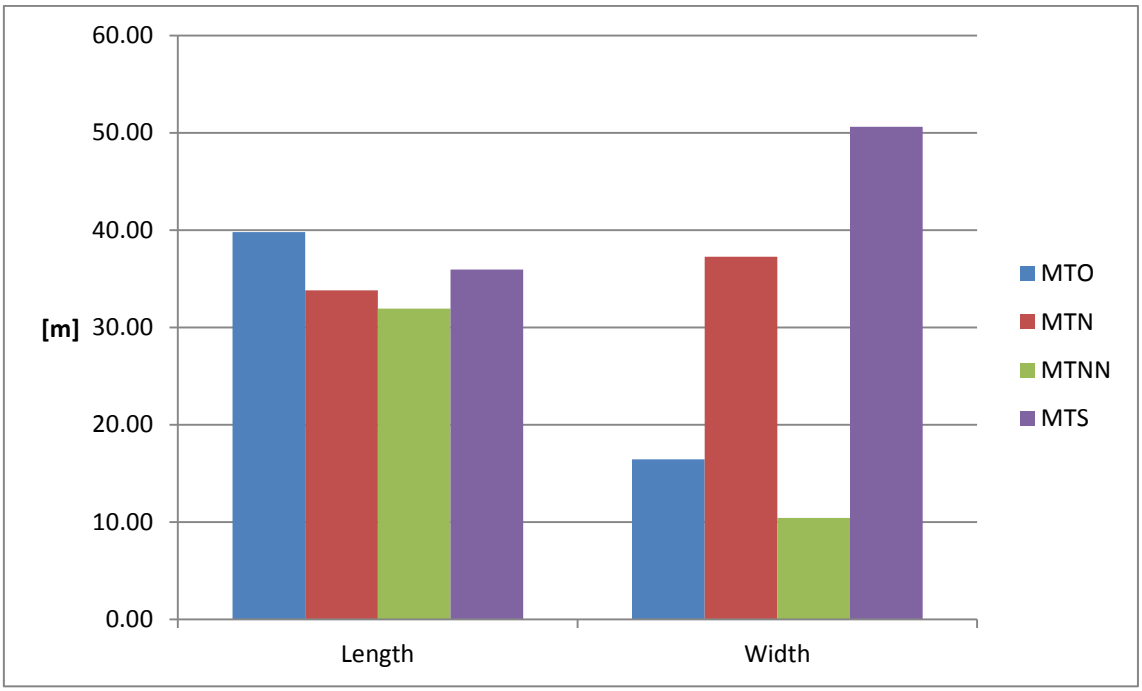

Figure 6-2. Mean absolute error (entire dataset).

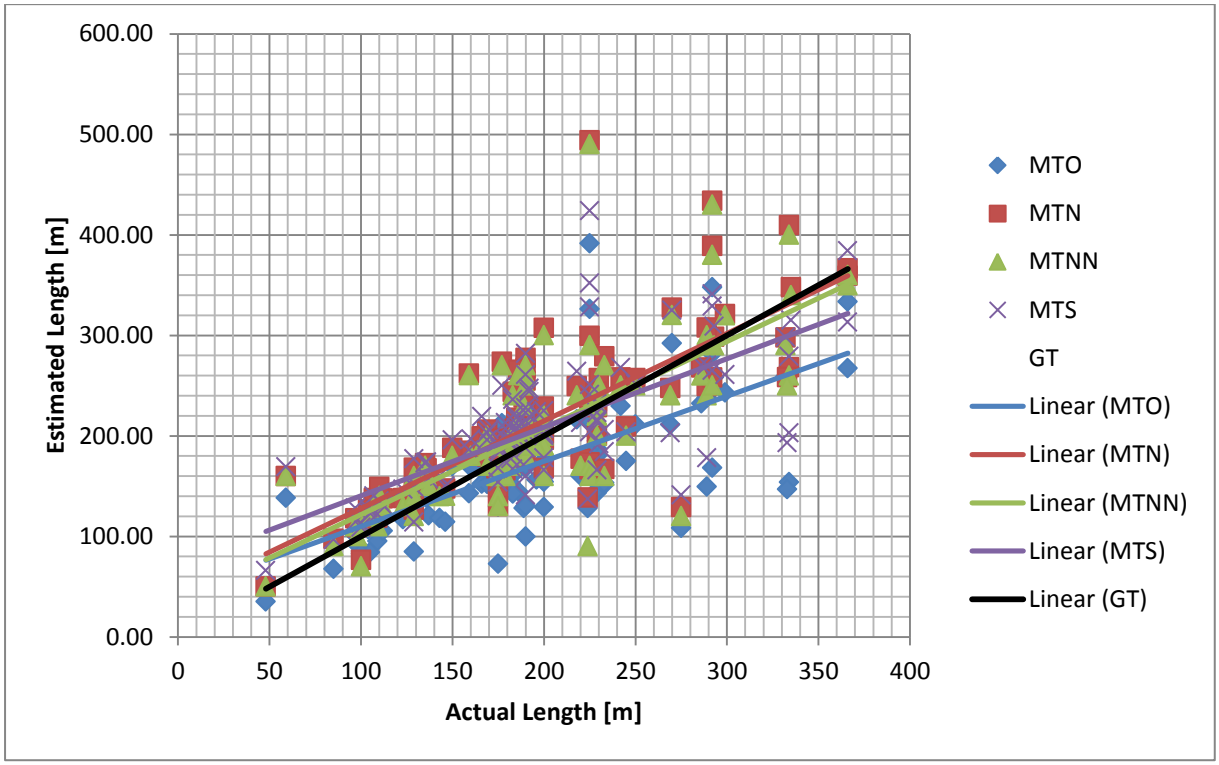

Figure 6-3. Length scatterplot (entire dataset).

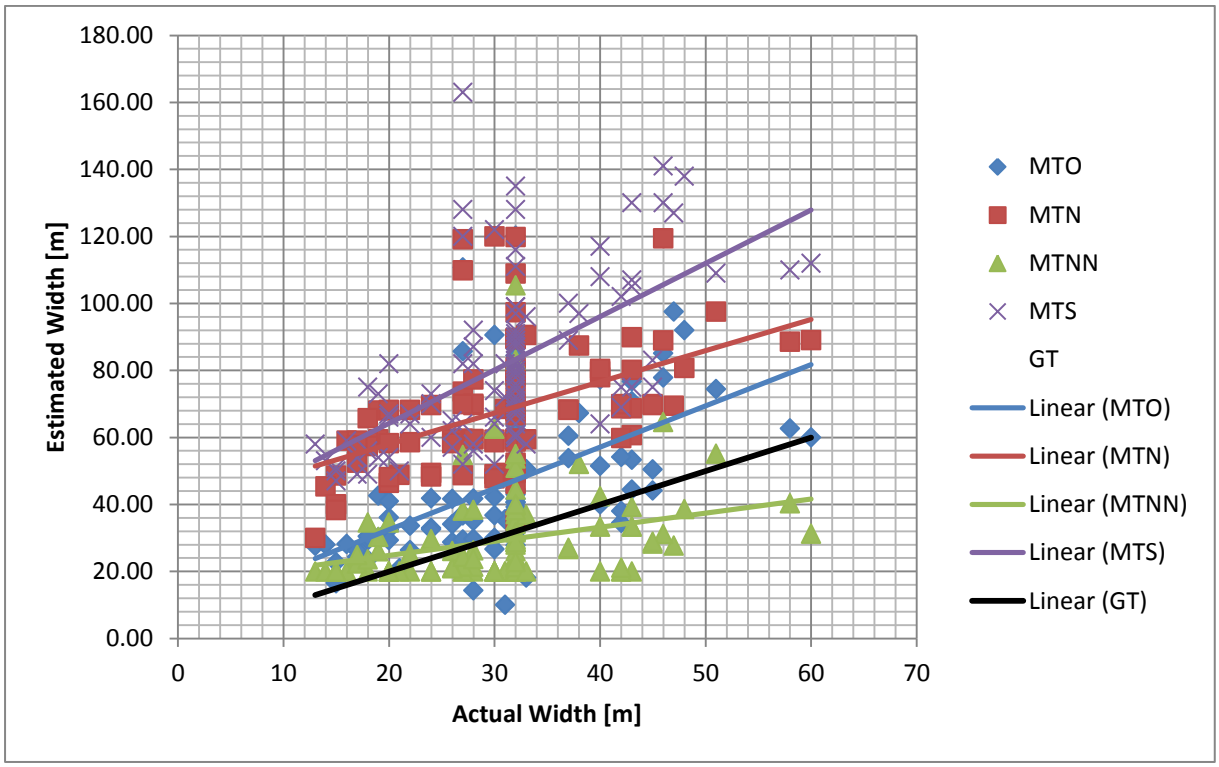

Figure 6-4. Width scatterplot (entire dataset).

6.2.1 Outliers

The overall results presented in the previous section well describe the behaviour of the four algorithms under test, as the errors reported in Table 6-2 are basically the same for every single image of the dataset. Nevertheless, in image 20150821T151840, the error in length estimation for MTN, MTNN and MTS was much worse than the one achieved over the entire dataset (see Table 6-3 and Figure 6-5). Except for MTO, the other 3 methods experience an increase of the mean error of about 15 pp. Width estimation is in line with the overall results and it seems not to be affected by particular issues.

It is therefore interesting to look in detail at the length errors, target by target, to find possible outliers. This is done in Figure 6-6, where we can immediately see that target 17 and 23 are characterized by two clear peaks significantly above the average error level.

To understand why the two outliers generated such high peaks, it is necessary to analyse their signatures.

According to AIS data, target 17 is a 59 x 15 m supply ship (Figure 6-7a) that should appear in the image as a small blob of a few pixels. Actually, Figure 6-7a shows that the target is smeared, as it features a bright "tail" due to the presence of the wake (the ship was moving at a speed of about 8 knots). As indicated by the white rectangle superimposed on the signature, the estimation algorithms did not separate the real target from the wake, and they output lengths with errors of about 100 m.

The same issue – even more noticeable – biased the results for target 23. It is a bulk carrier of 225 x 3 m, having a quite large and well defined SAR signature (Figure 6-7b). At the time of the image acquisition, the speed of the vessel was around 11 knots, thus a long wake was produced behind it. None of the 4 algorithms could distinguish the wake, thus the estimated lengths were about 100 m longer for MTO and MTS, and even 200 m longer for MTN and MTNN.

In order to prove that these two targets were the outlier worsening the performances, statistics have been re-calculated taking them out. The relative length errors dropped down to the expected values.

Table 6-3. Mean relative errors (entire dataset VS image 20150821T15184).

	Average Error	MTO	MTN	MTNN	MTS
Length (%)	Entire dataset	19.2	18.3	16.8	19.4
	image 20150821T151840	21.5	32.9	30.2	26.2
Width (%)	Entire dataset	55.2	132.6	33.6	174.2
	image 20150821T151840	42.1	134.8	31.0	169.0

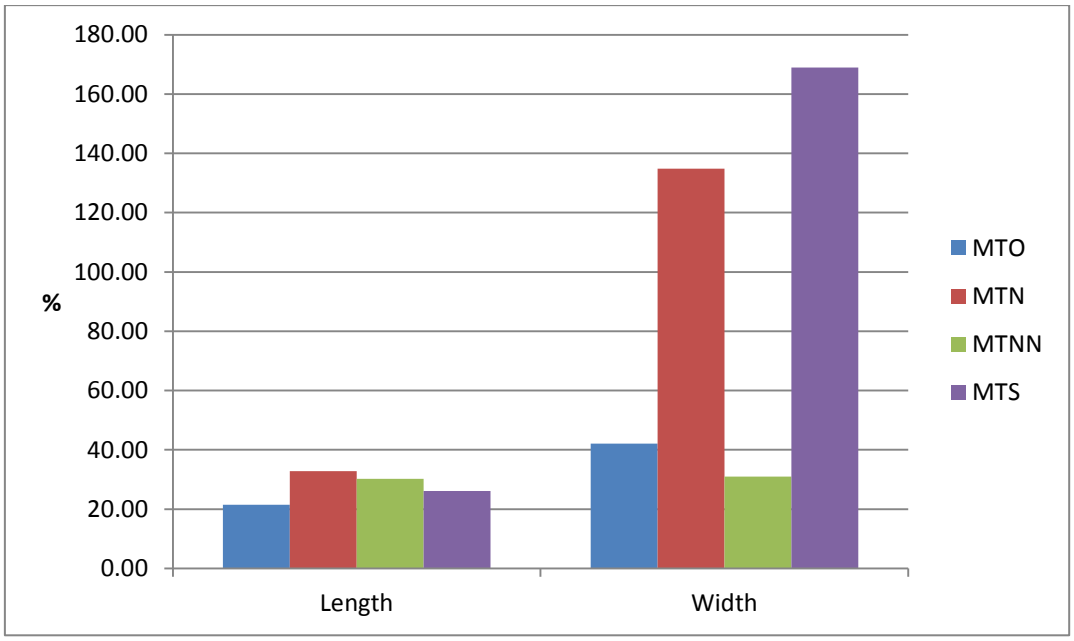

Figure 6-5. Mean relative error (image 20150821T151840).

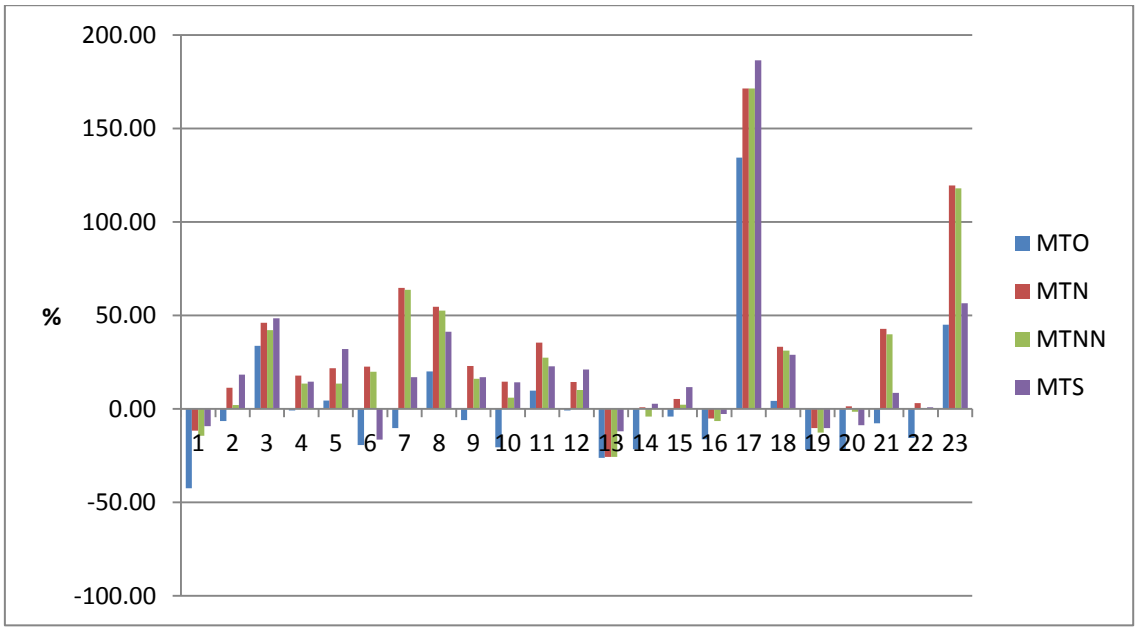

Figure 6-6. Single target relative error (image 20150821T15).

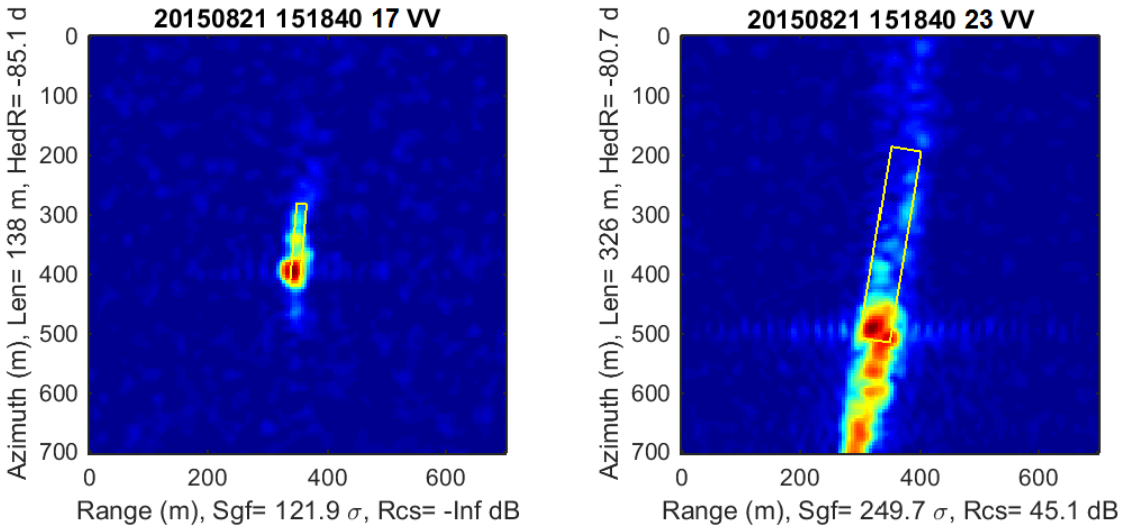

(a) (b)

Figure 6-7. SAR signatures of the outliers: (a) target 17; (b) target 23.

6.3 Conclusions

In this chapter, the benchmarking of Sentinel-1 for ship size estimation has been presented. The analysis has actually focused only on multi-looked detected IW GRDH products, with a resolution of about 20 m.

Four different estimation methods have been compared, which can all accomplish length reconstruction with good accuracy. At the given resolution, width estimation remains an unsolved issue, although a new algorithm recently developed could significantly reduce the errors.

In addition, it should be always taken into account that some outliers could bias the results due to undesired artefacts in SAR ship signatures. As already shown for other sensors [S2014], they cannot be always avoided and may reduce the overall accuracy.

7 Ship disambiguation in Sentinel-1 images

Ship classification, the ability to derive the ship class like tanker, cargo, fishing, passenger, etc., is very difficult to do with SAR images. Even if there is enough resolution to show many details, the SAR image of an object is quite different than the more familiar optical image, and in addition, any motions (such as pitching and rolling on the waves) result in serious blurring of the SAR image. For Sentinel-1, which is mostly operated in IW mode with the GRDH product at 20 m resolution, there is not even much detail in the images to start with.

In some cases, ship type can be narrowed down just by using the size. Such a narrowing down can be already useful in many scenarios; e.g. when looking for fishing ships, ships larger than 150 m can be excluded. Another situation where partial classification can be enough is in the correlation with positions from ship reporting data (AIS or LRIT). It often happens that two (or more) SAR ship detections occur in the vicinity of one reported ship position. The size and type of the reported ship are known. If the two SAR targets have different sizes, the correct one can be assigned to the reported position. However, if the two SAR targets have similar sizes, then the type has to be considered for the correlation.

In this chapter, we propose a method to disambiguate between cargoes and tankers when an association between a SAR signature and several AIS messages is needed and the size cannot disambiguate. The method proposes to analyse the visual SAR signature in order to classify it into one of the two vessel categories under discussion.

7.1 Dataset & ground truth

A dataset was created to perform ship size estimation, ship classification and ship type discrimination over Sentinel-1 images. The dataset was created using SUMO to detect the vessels present in the satellite images and correlating such information with available ship reporting data to ensure the accuracy of the ground truth. For further information on the analysed images refer to Section 6.1.

The satellite image analysis provided a total of 107 reliable targets. The targets presented a variety of AIS ship type codes mainly centred on cargoes and tankers. Their ship type codes vary between 80 to 89 and 90 to 99, meaning cargoes and tankers respectively, where the first digit indicates the ship type and the second digit refers to the transported goods. The selection of only cargoes and tankers reduced the ground truth to 98 ships, 67 cargoes and 31 tankers. Hence, the ground truth is composed of 98 ship samples, from which 30 samples are used in the training phase of the cargo/tanker disambiguation system to build the dictionary, using 15 samples from each vessel type. The complete ground truth was used to test the system.

The dataset described in Section 6.1 provides image chips of 140 by 140 pixels centred on the centroid of the detected vessel, independently of the vessel size. Figure 7-1 presents some examples of the analysed image chips. The image chips included in the ground truth have a limited resolution, defined by the satellite image resolution which in this dataset is of around 20 metres, and contain certain artefacts (e.g. sidelobes) inherent in SAR images.

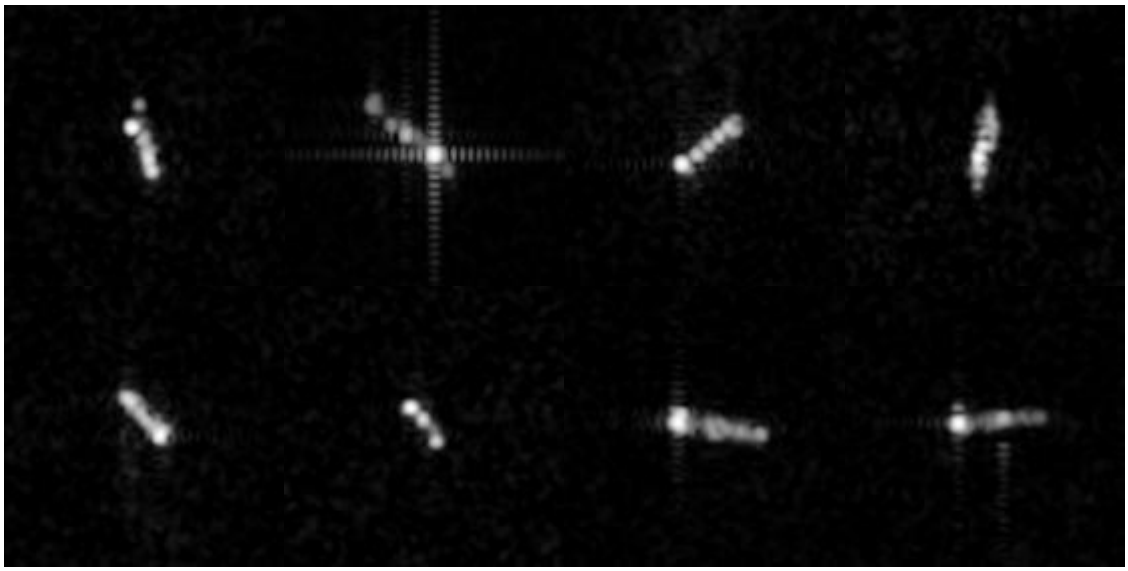

Figure 7-1. Image chips extracted from Sentinel-1 images using SUMO. All image chips have a size of 140x140 pixels. The upper row presents four samples of cargoes, while the lower row shows four samples of tankers. The second chip on top suffers from strong sidelobes.

7.2 Cargo/tanker disambiguation based on visual analytics

A method to distinguish between cargoes and tankers based on their visual appearance registered in the satellite SAR images is proposed. The method is based on the underlying idea that each target category shares common super-structures enabling vessel category classification. Figure 7-2 illustrates the framework of the proposed cargo/tanker classifier. The proposed method analyses image blobs extracted from satellite SAR images. SUMO, the in-house system for the satellite image processing and analysis, detects vessels and creates a list of possible targets for each satellite SAR image. For each detected vessel, an image chip is extracted.

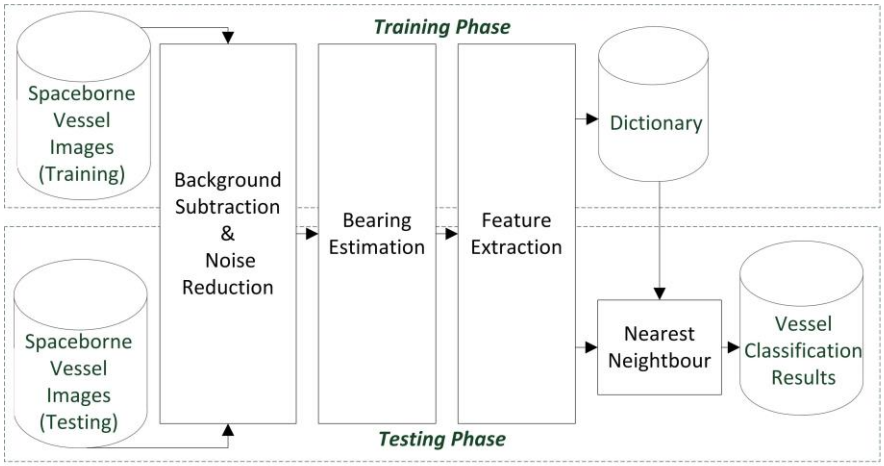

Figure 7-2. Cargo/tanker disambiguation framework. The system receives image chips, extracted using SUMO, and provides a classification level to each image based on its visual appearance.

The proposed cargo/tanker classification system is composed of two phases: training and testing. Whilst the main core of the approach is shared, each phase aims at a different objective. In the training phase, the training dataset is analysed to build up a dictionary, based on Bag of Visual Words [S2003]. In the testing phase, the extracted features are compared with the built-up

dictionary to determine the vessel class by detecting the nearest neighbour. The general structure of the system is based on the work presented in [F2015].

The framework presents several core stages where the method targets separating important visual information from noise in the image blobs. First, the image chips are pre-processed to remove the background and reduce the noise within the image. The image chip background covers a large percentage, hence, extracting visual features from the whole image chip would lead to results affected by external factors not related to the vessel. Second, to reduce the variability of the entries required to represent the vessel categories within the dictionary, all the vessel images must be invariant to rotation. In order to obtain rotation invariance, the bearing of the vessel is estimated. Then the images are rotated to a standardised position (0 degrees). Finally, the resulting image is analysed in the computer-vision feature extraction module, where different features are extracted to represent the vessel categories, considering their individual characteristics and the satellite SAR image restrictions.

Background subtraction & noise reduction

Maritime spaceborne images are typically affected by external factors modifying the state of the sea, i.e. wake, winds, currents, etc. Such effects on the images visual content hinder the distinction between the vessel and the background as well as having an impact over the extracted features. To limit their effects in the classification process, the proposed system processes the image chips to remove the background and reduce the noise. The proposed method is divided into three sequential stages:

1. A Sobel edge detector is applied performing a 2-dimensional spatial gradient measure over the image. As a result, the image high spatial frequency areas are emphasised and so the image edges are revealed.
2. Three morphology operations (dilate, fill and erode) are applied to join neighbouring elements and isolate independent elements.
3. Noise in the image affects the objects segmentation resulting in cases of over-segmentation. The connected components method is applied to group object sections and to disregard components composed by a negligible pixel amount.

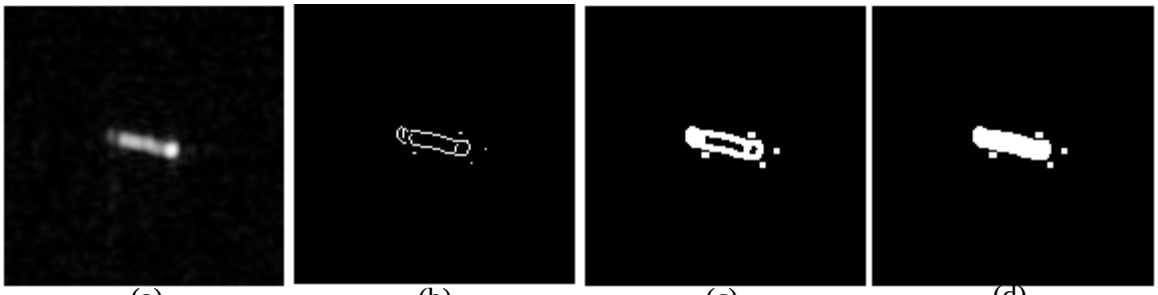

Figure 7-3. Background subtraction and noise reduction stages over a SAR image chip. (a) Original image chips. (b) Resulting image after applying a Sobel edge detector. (c) Obtained image after applying three morphological operators over the edge image (b). (d) Mask obtained applying connected components over image (c).

Bearing estimation

Prior to extracting the features and building the cargo and tanker signatures, the cargo and tanker image representations must be standardised, ensuring that even the use of features not invariant to rotation or translation would enable their representation. Two methods can tackle representability across all vessel positions: (i) to create a dictionary with different entries per vessel type *and*

orientation or (ii) to automatically estimate the bearing of the vessel and rotate all the image chips to an standardised position. The first option requires a large set of image chips to build the dictionary. The proposed algorithm estimates the bearing by calculating the orientation of the connected components computed in the previous stage. The method detects the major axis of an ellipse holding the segmented vessel and computes the angle between such axis and the x-axis.

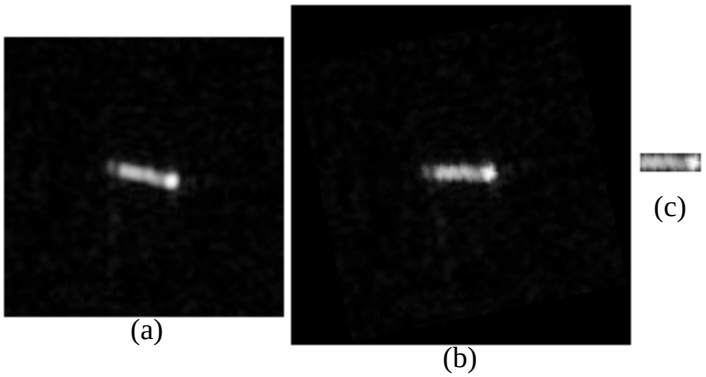

Figure 7-4. Bearing estimation. (a) Original image chip. (b) Rotated image, using the estimated length, width and heading measures. (c) Segmented image in the standard orientation position, the remaining blob size is 9 by 30 pixels.

Features extraction

In order to represent the vessels, the features selected need to fulfil several criteria. Generally, the features need to reduce the vessel type intra-variance while enhancing the inter-variance. More specifically considering the limitations of the analysed images (refer to Section 7.1), the features selected should enhance the characteristics still present in images with low resolution. Some features present in low resolution images are texture, orientation, edges or corners. Thus, we focused our study in these features, analysing their individual performance over the images, the correlation of the features across vessel types and their performance applying diverse feature fusion methods.

The analysed features include: Local Binary Patterns (LBP) and Histogram of Gradients (HOG).

- Local Binary Patters (LBP) presents a very efficient texture operator which analyses the neighbouring of each pixel, considering not only the local spatial patterns but also the grey scale contrast, whilst providing great discriminative power and computational simplicity [O1994, O1996]. The LBP features have been computed for the whole segmented image blob (refer to the results Table 7-1 as General LBP). However, considering the local distribution of texture characteristics within the image blob due to the vessel's super-structure, the image blob is decomposed in three segments being bow, middle and stern. Such decomposition consists of the division of the segmented vessel into three equal parts computed from the estimation of the geometrical characteristics and ratio of the vessel. Such decomposition is used to compute the LBPs in each vessel part, extracting the texture characteristics more distinguishable in each segment (refer to the results Table 7-1 as Bow LBP, Middle LBP and Stern LBP, respectively).

- Histogram of Oriented Gradients (HOG) intends to represent the object's appearance and shape calculating the distribution of intensity gradients or edge directions. HOG divides the image blob in equal-size cells and computes a histogram of gradient directions for each cell, concatenating the results to obtain the feature descriptor. HOG is a powerful and discriminative descriptor presenting invariance to geometric and photometric transformations but not to orientation (the reason why bearing estimation is computed beforehand) [D2005]. The HOG foundation is based

on the computation of gradients, based on derivatives. Hence, HOG intends to detect the object edges, considering the object shape and edges are the key factor towards their detection and modelling. HOG has been applied to the segmented object (refer to the results Table 7-1 as General HOG).

Testing phase

The training phase ends with the construction of a dictionary, based on Bag of Visual Words concept, composed of a representative set of feature samples from each vessel category, in our case, cargoes and tankers. The testing phase, on the other hand, follows the same core stages to extract the same features and build the visual words that are later compared to all the entries within the dictionary. A classifier applying nearest neighbours based on the Euclidean distance is proposed.

7.3 Experimental results

The presented approach proposes a cargo/tanker disambiguation based on visual analytics, building the approach on robust features capable to represent vessels in low resolution satellite images. Table 7-1 presents the individual results of the feature analysis. The table presents the accuracy of the features measured as the true positive detection rate in a classification problem.

Table 7-1. Features individual classification results for cargo/tanker disambiguation. All figures in %.

	General LBP	LBP Bow	LBP Middle	LBP Stern	HOG
Cargo	59.7	70.2	67.2	77.6	68.7
Tanker	71.0	77.4	71.0	64.5	67.7
Across ship types	63.3	72.5	68.4	73.5	68.4

The results reveal two features exceeding the achieved average accuracy, being LBP Bow and LBP Stern. On the other hand, the General LBP presents the lower accuracy across ship types. The HOG feature presents an stable and average performance across ship types.

To analyse the features intra-class representability and inter-class variability, we computed the correlation across image blobs in the Training dataset for every feature extracted and the Pearson correlation coefficient. Establishing the requirements of a confidence level of over 90 % and a Pearson coefficient over 40 %, certain features show statistical significance and reveal possibilities to predict amongst vessel types, including General LBP, LBP Middle and HOG.

The initial hypothesis established the need of using features with high intra-class similarity, and thus, high correlation amongst the images selected for the training dataset for each vessels class. However, cargoes and tankers differ largely according to the transported material, the carried containers and many other factors (refer to Figure 7-5). Thus, to build a strong dictionary, Bag of Visual Words, we need to include large variability within the vessel classes, enabling the identification of cargoes or tankers despite their difference in the appearance. Moreover, in Figure 7-1, we present some examples of cargoes and tankers, presenting their appearance on Sentinel-1 images and their dissimilarity across signatures, despite their belonging to the same vessel category.

Figure 7-5. Tanker appearance variability.

Despite the great limitations presented by Sentinel-1 images for vessel classification, individual features present promising results for ship categorisation according to their visual characteristics, based on the analysis of some features whose representability is limited but not eliminated by the low resolution of the images. However, the complementarity of the analysed features builds the hypothesis that a feature combination could benefit the performance results. Therefore, the next section looks at feature combinations.

7.4 Multi-feature cargo/tanker disambiguation

Generally, feature fusion can happen either at feature level or at decision level. Feature level fusion merges features to create a unique descriptor that would pass as an individual to the classifier, requiring the merge of features living in the same feature space. On the other hand, decision level fusion combines the individual features' classification results, ensuring the features independence and enabling the computation of a certainty probability which indicates the certainty level of the classification [F2014]. Considering the pros and cons of feature level and decision level feature combinations, we propose to focus on decision-level fusion due to the features' different characters and element significance (refer to Figure 7-6 for the decision-level classification schema).

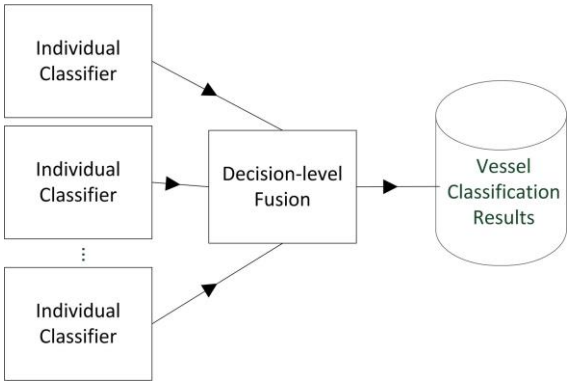

Figure 7-6. Decision-level classification schema.

Several feature combinations and fusion techniques have been tested for our objective. First, considering the results obtained from the individual features evaluation, we proposed major voting as fusion technique at decision-level. Using major voting we proposed three combinations, including:

- Combination 1: equally-weighted linear combination of General LBP, LBP Bow, LBP Middle and LBP Stern.
- Combination 2: equally-weighted linear combination of General LBP and HOG.
- Combination 3: equally-weighted linear combination of LBP Bow, LBP Middle, LBP Stern and HOG.

Table 7-2 shows the results. The results reveal an improvement in the classification performance, reaching 81.6 % accuracy across ship types and even 88.1 % and 67.7 % classification accuracy for cargo and tanker, respectively. From the results in Table 7-2, we can conclude that Combination 3 provides a stronger descriptor and stronger ship disambiguation capabilities. Moreover, a

comparison between the combined results and the individual feature results reveal a performance improvement by 8.1 percentage points (from the best individual feature at 73.5 % to the best feature combination at 81.6 %).

Table 7-2. Decision-level feature fusion classification results for cargo/tanker disambiguation. Figures in %.

	Combination 1	**Combination 2**	**Combination 3**
Cargo	76.1	82.1	88.1
Tanker	71.0	58.1	67.7
Across Ship Types	74.5	74.5	81.6

7.5 Results analysis and conclusions

In this chapter, a method to disambiguate between cargoes and tankers based on visual analytics has been presented. The proposed framework analysed image chips extracted from Sentinel-1 IW GRDH images with a resolution of about 20 m.

The experimental results are based on the visual analysis of 98 image chips, from which 67 represented cargoes and 31 represented tankers. The method proposed to standardise the image chips to a particular orientation and segment the vessels from the image chips in an attempt to reduce the background and noise effects on the classification performance. After that, two features were proposed to represent the vessel types under study. In the proposed method, vessel structure was seen as the main remaining characteristic on low resolution satellite images. Thus, the features selected, LBP and HOG, envisaged to extract patterns representing the vessel structure. Their individual performances ranging from 63.3 % to 73.5 % across vessel types are promising results, considering the size of the segmented images (refer to Figure 7-4(c) for an example on a segmented vessel image and its size). Subsequently, a decision-level feature fusion is proposed to perform the cargo/tanker disambiguation based on multiple features. Combining features not only improves the performance – increasing up to 81.6 % –, but also provides a classification certainty, establishing the confidence on the classification label.

SAR ship signatures can have serious artefacts. These can impact on the proposed analysis. This may cause errors in the bearing estimation and in the segmentation outcomes as well as introduce outliers in the classification stage, affecting the overall accuracy. Some examples of the effects of undesired artefacts are shown in Figure 7-7.

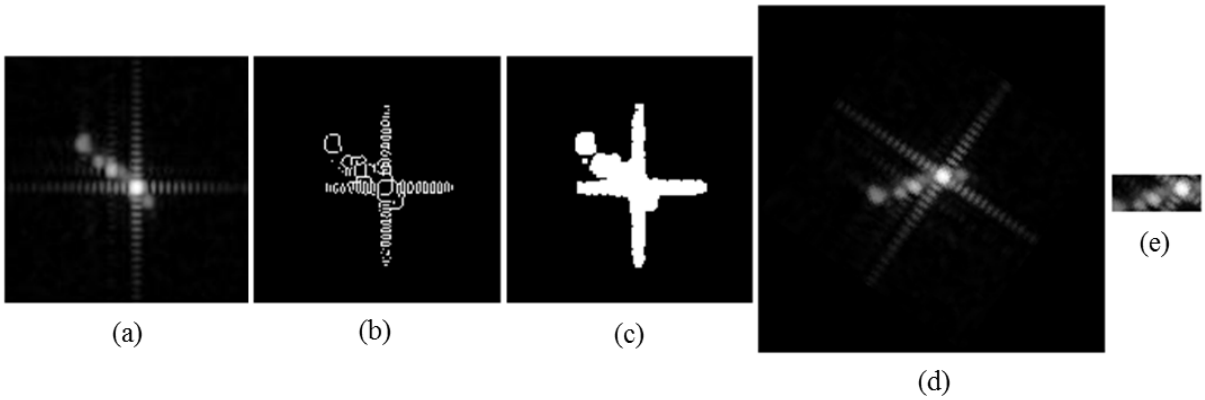

Figure 7-7. Vessel image chip affected by undesired artefacts, common in SAR ship signatures. (a) Original image affected by sidelobes. (b) Detected edges. (c) Mask obtained after applying detected components. (d) Image rotated. It should be rotated to the horizontal line, however, the sidelobes affected the bearing estimation, resulting in the vessel signature leaning around 30 degrees over the horizontal line. (e) Segmented image. The resulting segmented image is over-segmented due to the error on the bearing estimation, such error accumulates across the system.

8 Other issues

8.1 Duplicated ships in Sentinel-1 TOPS SLC images

The most useful product type for bulk processing of SAR images for maritime surveillance is Ground Range Detected (GRD). This is the product type that has been used in the Western Indian Ocean and Med Sea campaigns reported in this document. Products of the type Single Look Complex (SLC) are less suitable for bulk processing (among other things because of their much larger data size than GRD products) and are still largely seen as an experimental type. (For other applications such as interferometry, they are the standard product type.) ESA very rarely releases SLC products in the EW and SM modes, but it now systematically releases most of the IW acquisitions as both GRD and SLC products. A few of these IW SLC products have been studied in the course of the last year.

The Sentinel-1 IW and EW modes use an antenna scanning pattern called Terrain Observation by Progressive Scans (TOPS). TOPS sends a short burst of radar pulses in one of the mode's sub-swaths, then switches to the next sub-swath to send a similar burst, and so on until the last sub-swath. Then the antenna is aimed at the first sub-swath again and the cycle repeats. In the case of IW image the burst cycle lasts 2.75 seconds. There is also some overlapping between consecutive bursts in a sub-swath. This means that the overlapping area on the ground/sea will be observed twice by Sentinel-1 in TOPS mode, with a time different between observations of around 2 seconds and with slightly different squint angle. This time diversity can be used to study time-variant objects (like moving ships) in SLC products, since in this product type the bursts are presented in their entirety. In the GRD products the pixels shown in the overlapping areas are selected from one of the bursts only (an operation called de-bursting), thus eliminating the "dual" observation. The overlapping area between consecutive bursts in the IW mode has been approximately 2 km since the start of the operational phase.

A number of ships have been detected in the overlapping areas of IW SLC images. The majority of these ships present very similar signatures in the two consecutive bursts, but a few of them show markedly different features. Figure 8-1 (left) shows a section of an SLC image that includes parts of two consecutive bursts and the black zero padding area that separates the bursts. There are four ship signatures in the figure, but the two signatures on the left correspond to the same ship, which is located in the burst overlapping area. The observed heading of the signatures of this ship are very different: in the upper burst the ship seems to be heading up (the stern is on the left, the bow is on the right), whereas in the lower burst it seems to be heading down. A large ship cannot change its heading so abruptly in such a short period of time (2 seconds), so this is probably a case of distorted signatures due to ship motion. This effect has been seen in other SARs ([O2002] estimates that ship pitching is the most significant contributing factor to the distortion), but the "dual" observation of Sentinel-1 makes it very apparent. The consequence of this motion-induced distortion is that it is not possible to ascertain what the real heading of the ship was just by looking at the SAR image: it could be the one observed in the upper signature, the one observed in the lower signature, anything in between, or even a value outside that range. Estimates of the ship's length will also be affected by the distortion.

Figure 8-1 (right) displays the same area in the equivalent GRD image. The de-bursting operation has produced an image without gaps that shows the same three ships. It has to be noted that the resolution in the range direction (i.e. horizontally in Figure 8-1) in the SLC images is much higher than the resolution in the azimuth direction (i.e. vertically), thus producing ship signatures that are elongated horizontally. On the other hand, GRD products have the same resolution in range and azimuth, so the signatures appear horizontally shortened with respect to their SLC counterparts. In the GRD image the signature of the ship on the left has been selected from the upper burst of the SLC product. This is so because the signature in the upper burst is further away from the black zero padding area than the signature in the lower burst. AIS has been used to corroborate that the real

heading of this ship at that time was similar to the heading seen of the other two ships in the image. So in this case, a SAR-based heading estimate would be very inaccurate.

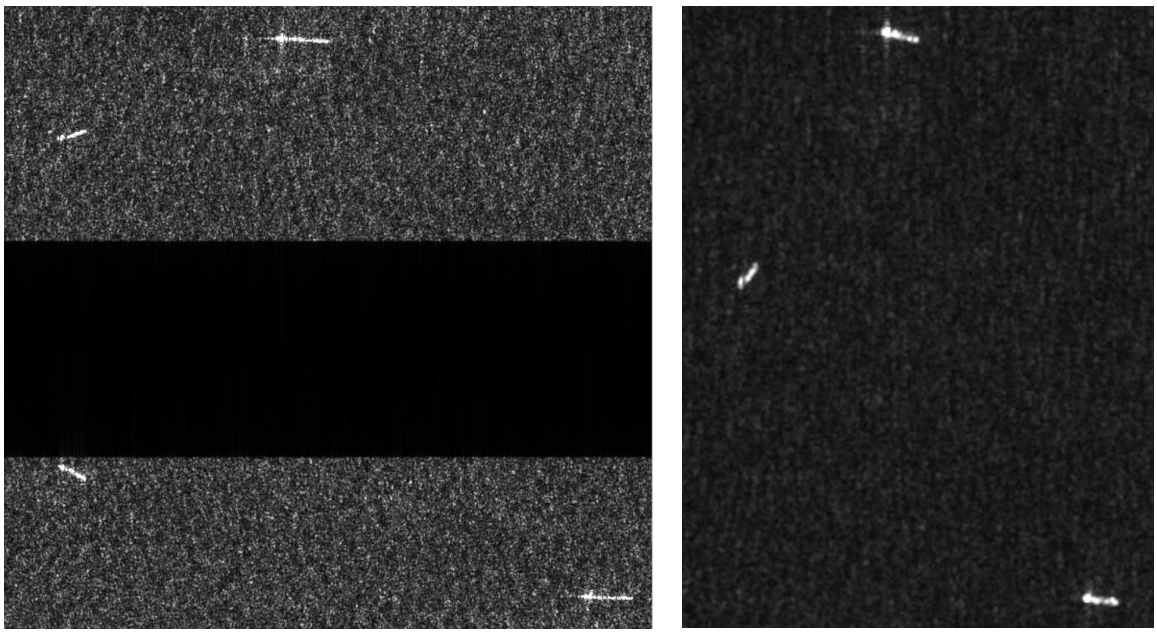

Figure 8-1. Sentinel-1 IW image with a ship located in the burst overlapping area (ship on the left of the images). Left is SLC product, right is GRD product. In the SLC product the ship is observed with noticeably different headings. In the GRD product one of the signatures in the overlapping area has been selected. AIS data indicate that the real heading of the ship was similar to the heading seen of the two ships on the right of the images.
S1A_IW_SLC__1SDV_20150227T061013_20150227T061045_004805_005F94_6248
S1A_IW_GRDH_1SDV_20150227T061015_20150227T061045_004805_005F94_DFCC

Motion-induced ship signature distortion is also clearly seen in Figure 8-2. The ship wake (the darker line extending behind the ship and reaching the far left border of the figure) and AIS data contradict the ship heading observed in the SAR image. The SLC product is at the top, the equivalent GRD product is below.

The squint angle with which a target is illuminated by the radar can significantly alter the SAR signature of the target. An example is seen in Figure 8-3 where a container ship in the burst overlapping area shows a very saturated signature with visible side lobes in the lower burst, and a clearer signature in the upper burst. This change in the signature may be due to a slight change in the squint angle (as a result of the TOPS scanning pattern), perpendicular to the container stacks in the lower signature, slightly off the perpendicular angle in the upper signature.

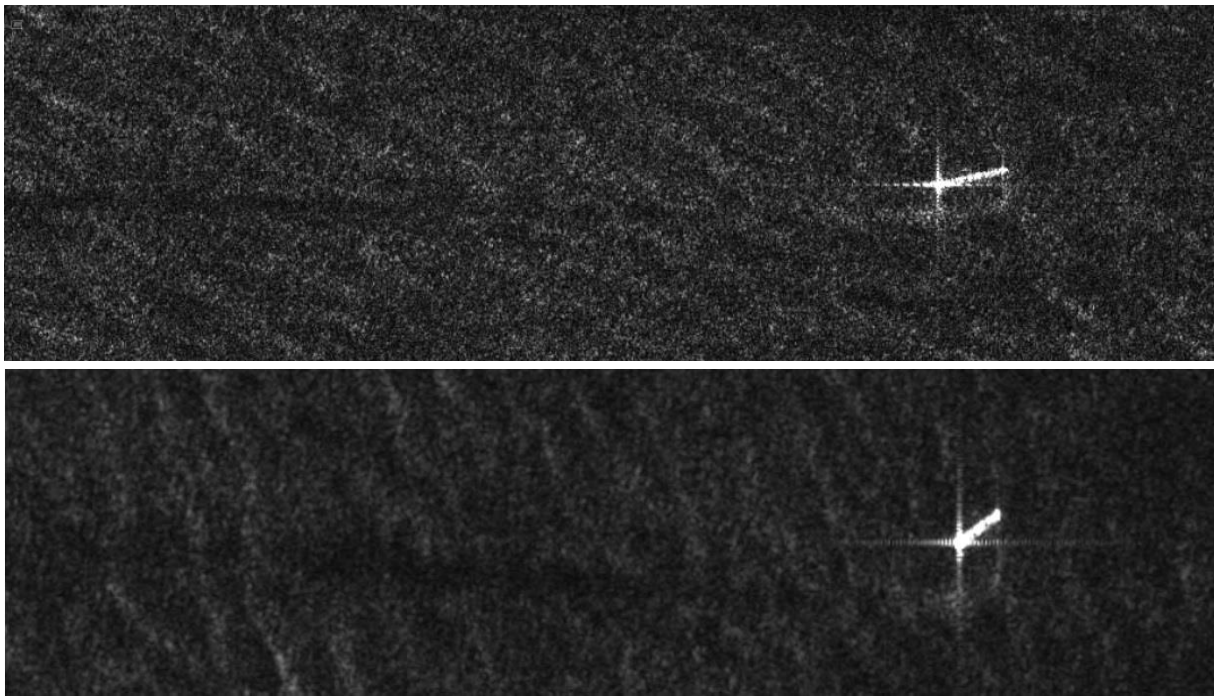

Figure 8-2. Sentinel-1 IW image where the ship signature and the ship wake (darker line behind the ship) vary significantly. AIS data corroborates the wake-based heading estimation. Top is the SLC product, bottom is the GRD one.
S1A_IW_SLC__1SDV_20150307T064121_20150307T064151_004922_00625E_7CDD
S1A_IW_GRDH_1SDV_20150307T064121_20150307T064150_004922_00625E_7329

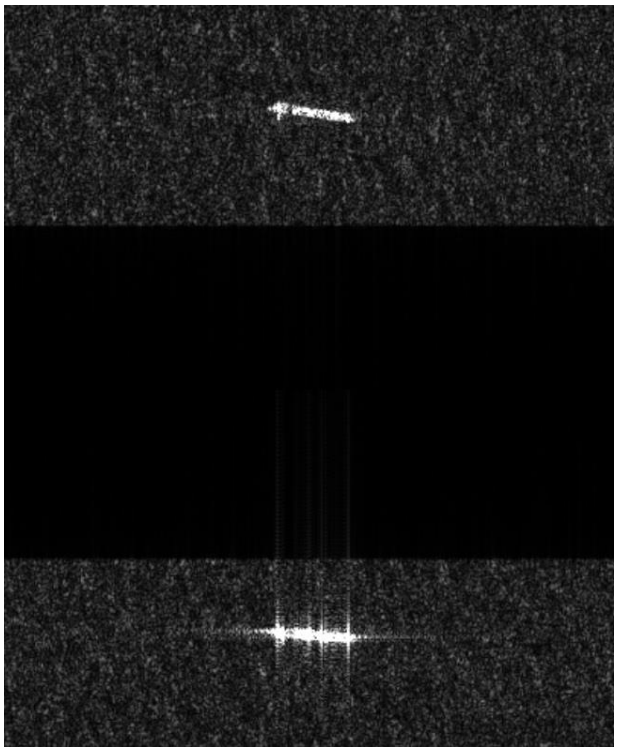

Figure 8-3. Ship in the burst overlapping area with markedly different saturation levels, maybe due to different radar squint angles.
S1A_IW_SLC__1SDH_20150129T055746_20150129T055805_004382_005593_EAFA

8.2 Analysis of range ambiguities affecting maritime images

Range ambiguities are ghost signatures of strong scatterers that appear displaced in the range direction from the originating scatterers. They are a cause of concern for maritime surveillance because they can give rise to false alarms in the ship detection process. In the case of Sentinel-1 IW mode, the distance between the originating target and the range ambiguity can be from approximately 110 to 220 km, depending on a number of observation and sensor parameters. The method being currently employed by JRC to ascertain whether a detection is a range ambiguity or not is by making use of the repeat acquisition operations of Sentinel-1 and the fact that ambiguities will show recurrence in a location for a given set of sensor and geometric parameters (mode, beam, and orbit track) and not for other sets of parameters [SG2015].

Another method to discriminate range ambiguities is to look for the existence of strong scatterers at the theoretical ambiguity distance from the detection; if a strong scatterer exists, the detection is likely to be a range ambiguity; otherwise it is not an ambiguity but a real object. This method is equivalent to the approach that is used to discriminate azimuth ambiguities. The large range ambiguity distance (compared with a typical azimuth ambiguity distance of around 5 km) means that the originating target is often outside the image where the range ambiguity is seen and, therefore, it has to be looked for in another image. Even if the two images were acquired in quasi-parallel neighbouring tracks, this adds uncertainty to the discrimination process since the originating target will be illuminated by the radar in the two scenes under slightly different incidence and squint angles, which could affect the strength of the originating target signature or of the ambiguity signature. Moreover, the acquisition times of the two images will be days apart, making the discrimination of ambiguities caused by time-variant structures (e.g. container terminal in ports, a very common source of ambiguities) difficult. For these reasons, the method based on recurrence is currently preferred. Nonetheless, an analysis of a few images containing range ambiguities and their originating targets has been carried out.

Figure 8-4 presents an example of a sea area with several range ambiguities (top), and the area in a quasi-parallel image containing the originating targets (middle). The images have been stretched and contrast-adjusted to allow for an easier visual comparison. A good (but not perfect) match between the ambiguities and the originating targets is observable. Figure 8-4 (bottom) also shows the footprints of the two images, indicating the position of the ambiguities (red star in the red image frame) and of the originating targets (green star in the green image frame). The images were acquired with a time difference of 7 days, and it is estimated that the originating targets were observed with an incidence angle difference of less than 10 degrees.

Figure 8-5 shows an example where range ambiguities and originating targets are present in the same image. The match between ambiguities and targets seems to be better than in Figure 8-4.

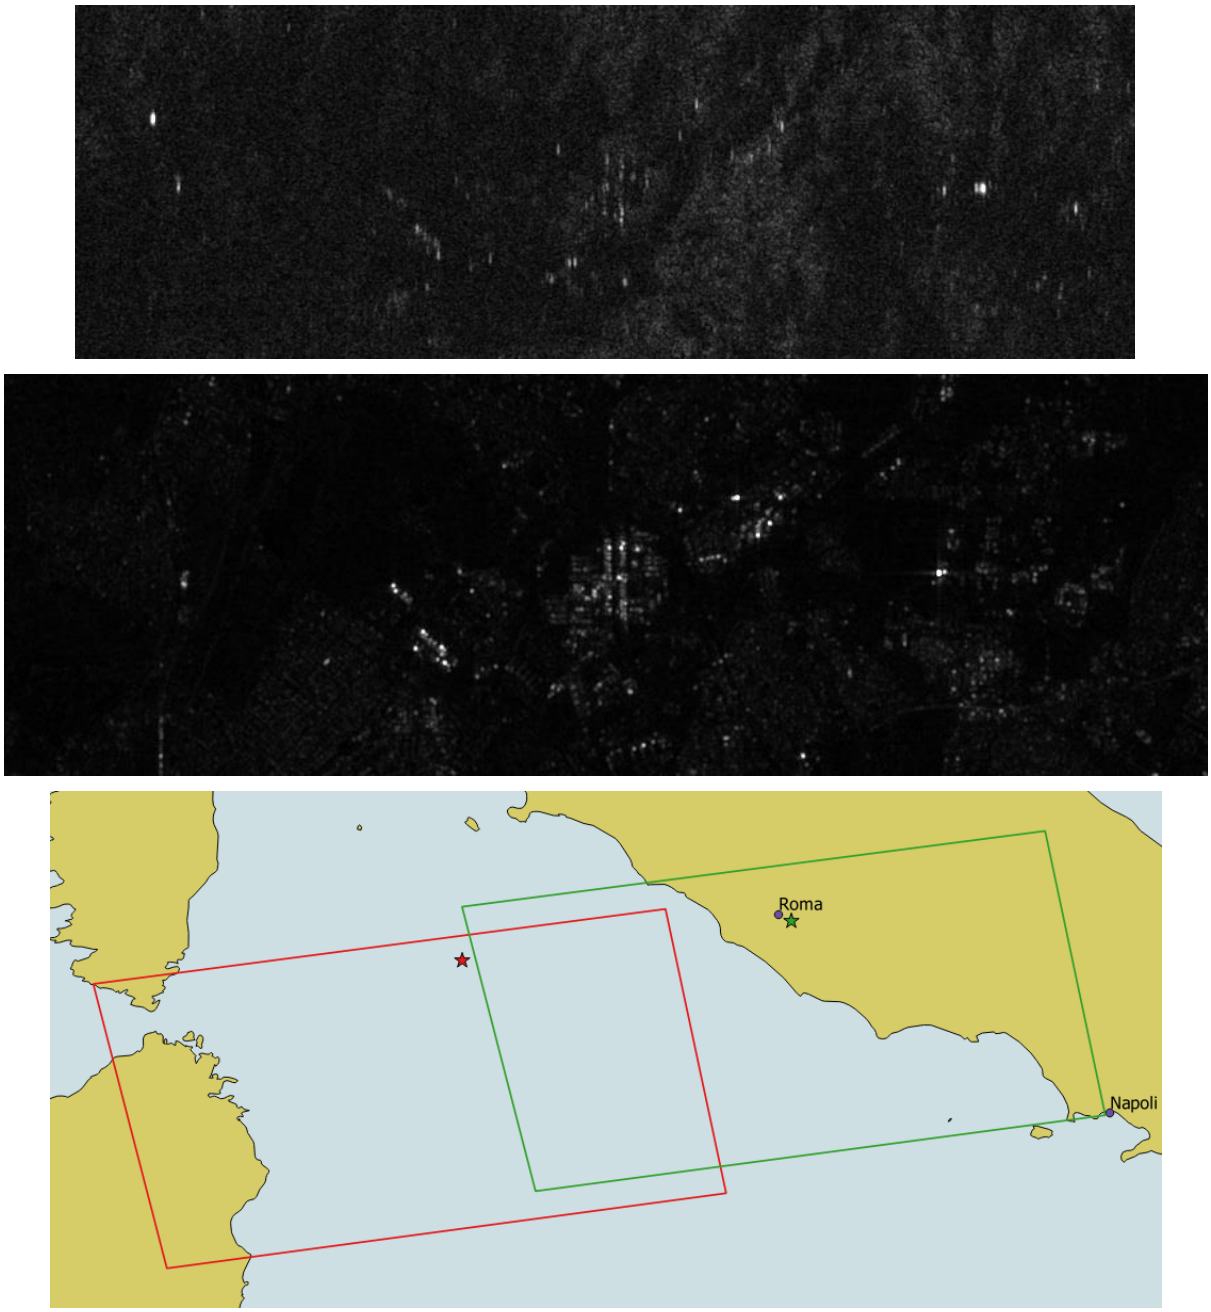

Figure 8-4. Top: range ambiguities observed in a Sentinel-1 IW image acquired on 23-03-2015 at 17:13h (UTC). Middle: originating targets observed in an IW image acquired on 30-03-2015 at 17:05h (UTC). Image intensity has been lowered to allow for visual comparison with the ambiguities. Bottom: map indicating the frame of the image containing the range ambiguities (red box), the frame of the image containing the originating targets (green box), the location of the ambiguities (red star) and the location of the originating targets (green star, near Rome). S1A_IW_GRDH_1SDV_20150323T171312_20150323T171337_005162_00681A_B944 S1A_IW_GRDH_1SDV_20150330T170506_20150330T170531_005264_006A6C_37FE

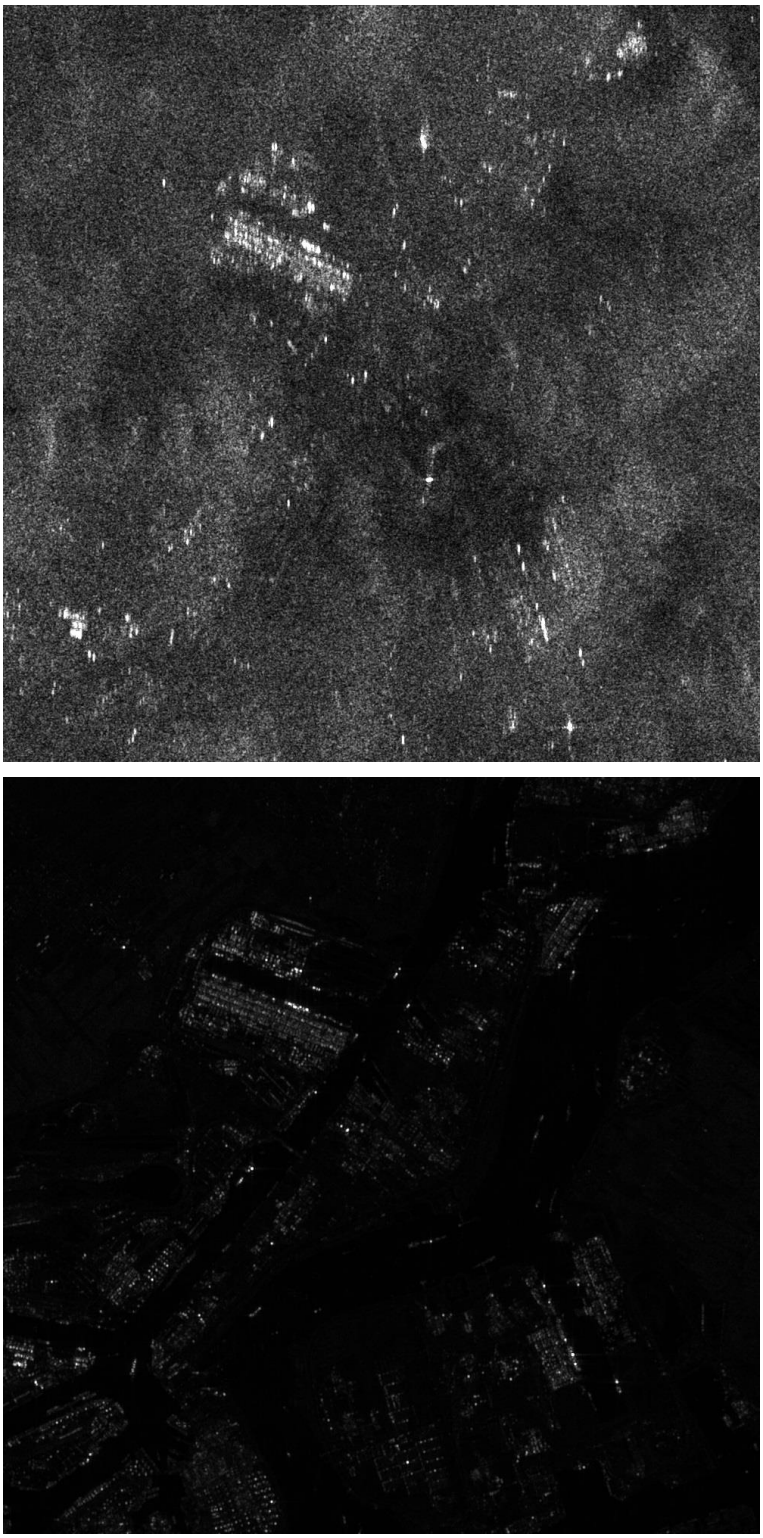

Figure 8-5. Range ambiguities (top) and originating targets in the port of Antwerpen (bottom), observed in the same image. Image intensity of the image containing the originating targets has been lowered to allow for visual comparison with the image containing the ambiguities. S1A_IW_GRDH_1SDV_20150210T055804_20150210T055829_004557_0059A3_8CB0

9 Conclusions

It is now possible to process large amounts of Sentinel-1 images that are being acquired over the sea for ship detection, fully automatically without human operator intervention on the separate images, and derive meaningful ship detections that are not plagued by false alarms. This capacity comprises:

a) Downloading of the images from the ESA rolling archive to JRC;

b) Running SUMO, the ship detector, on the images;

c) Correlating the ship detections from the Sentinel-1 images with known ships from the ship reporting systems (in the Blue Hub);

d) Checking for recurrent targets in repeat-pass images and for targets in overlap zones.

Processing times for all this are not excessive. These steps took in total 14 minutes per image on average in the set-up that was used, that is without parallelisation (except for the use of 2-4 cores in the available computers). These times are also not optimised since, due to IT architecture limitations, the images were stored on a different computer than where they were processed, introducing unnecessary network transfer times. Therefore, it is estimated that on an improved architecture and with the four processes (a) to (d) running in parallel, 400 images per day can be processed. This number can be compared with the Sentinel-1 production rate, which was e.g. during October 2015, 469 images over the Med Sea and 3,786 images over the Arctic, so that it would take just over 1 day to process one month worth of Med Sea images and 9.5 days for one month worth of Arctic images.

The scale of the processing executed for this report was limited by storage capacity. Being a development, it is desirable to store all the data once downloaded for repeated re-processing. But of course if an operational application is needed (none are planned at JRC), images need not be kept after processing; in fact, such processing would better be done at the image archive, obviating the need for the download step.

The routine data acquisition of Sentinel-1 has made it possible to analyse long time series, of many months to more than one year. In combination with the ship self-reporting data (AIS and LRIT) this has given new insights and understanding into the human activities at sea. In the Arctic in the presence of sea ice, it is really the ship reporting data that provides the information on the ship traffic, as the ship signatures in the SAR images drown in the background of the sea ice. However, in that case Sentinel-1 provides the information on the sea ice extent and dynamics that help to understand the shipping patterns. In areas where Sentinel-1 coverage is frequent, such as the Med Sea, the long-term image collection enables to map out ship traffic patterns, of reporting and non-reporting ships alike[1]. In remote areas that are not so well surveyed but where there is nonetheless an interest in the ship traffic, such as the Western Indian Ocean, satellite AIS is a very powerful tool to obtain maritime awareness concerning the larger ships (as AIS is mandatory for the ships of 300 Gross Tonnes and up). In order to also get an insight into the presence and distribution of the non-reporting ship traffic, which is the smaller ships but may also be larger ships of which no AIS reports are received, the combination of Sentinel-1 images with the ship reporting data provides the right tool.

For the purposes of maritime surveillance, it is desirable to obtain as much information as possible about the ships detected in the Sentinel-1 images, besides their geographic position. However, the Sentinel-1 images come at a resolution of 20 m (IW mode, GRDH product) and 50 m (EW mode, GRDH product). This means that size estimates are at best limited to such an accuracy, and in reality the accuracy is worse as the SAR signatures are subject to distortions and blurring caused by the

[1] As a caveat, it is know from past studies that the fixed time of day that the images are acquired gives rise to sampling effects on ship traffic that also observe a fixed daily schedule such as ferry routes; an effect that may however also occur with the reporting data, as AIS satellites also have fixed times of overpass and LRIT has a fixed 6-hourly reporting schedule. These effects were not explored with the data in this report.

motions of the ship and the surrounding water. Several automatic algorithms for size estimation were tested, and it was found that the best one gives a mean error of 32 m for the length estimate in IW GRDH images.

As a step in the direction of deriving the ship type (ship classification), an algorithm to distinguish between the types tanker and cargo was developed, based on the combination of two texture measures that are applied to a SAR ship signature that is segmented from the background. The results show a classification accuracy over these two classes of 82 %. Even though the problem (tanker-cargo discrimination) is a small subset of a full ship classification, this accuracy is still high considering the limited amount of structure that is present in the 20 m resolution signatures.

Concerning future developments, many improvements are still possible. Besides optimising the IT architecture as discussed at the beginning of this chapter, the SUMO software should be further improved on a few accounts. Some improvements with a large impact that are feasible to achieve in the near future include:

- The use of better coastlines by making use of additional layers in the OSM dataset that cover tidal areas and other coastal features;
- Better automatic rejection of small targets that are most likely to be false alarms, which would allow to lower the co-pol detection thresholds that are now used as a crude way to avoid excessive false alarms;
- Automatic downloading and using of ice edge maps to enable ship detection in the Arctic near the ice edge (preliminary tests have already been done for the downloading using the ERDDAP protocol).

Beyond SUMO, the recurrent targets framework that was started in 2015 should be further developed, as it is eminently suitable to deal with the repeat pass nature of the Sentinel-1 images.

Concerning the detection of small targets, this is a point of high interest but it is inherently difficult to search for small targets over wide areas in SAR images. Under most sea conditions, sea clutter gives too many false alarms, and even under optimal conditions in the absence of false alarms, one can expect many small boats to be present that are legitimate, without good indications about which ones could be problematic. The Sentinel-1 IW mode is, at 20 m resolution, not the best for small boat detection (and the EW mode even less so). Sentinel-1's SM mode (in its GRDF type, with 9 m resolution) would be more promising for small boat detection, but it is hardly ever provided.

It is hoped that the results in this report can be helpful in the development of the Copernicus maritime services.

References

A2010. Marlene Alvarez Alvarez, Juan Cicuendez Pérez, Tulay Çokacar, "JRC Vessel Detection Tool, SAR based (SUMO 4.1): Model description", JRC Scientific and Technical Report, JRC63144

AA2015. A. Alessandrini, provided on 26 Nov 2015. Background map made with Natural Earth. Free vector and raster map data @ naturalearthdata.com

BH2015. JRC, Web page of Blue Hub, https://bluehub.jrc.ec.europa.eu/, accessed 27 Nov 2015

C2015a. European Commission, http://www.copernicus.eu/main/copernicus-brief, accessed 30 Nov 2015

C2015b. European Commission, "Copernicus, Europe's eye on Earth", Publications Office of the European Union, 2015, ISBN 978-92-79-45666-4, doi 10.2873/93104

D2005. N. Dalal and B. Triggs. "Histograms of Oriented Gradients for Human Detection", *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Vol. 1, pp. 886–893

E2013. European Space Agency Sentinel-1 Team, "Sentinel-1 User Handbook", GMES-S1OP-EOPG-TN-13-0001, 1 September 2013

E2014. European Space Agency Sentinel-1 Team, "Sentinel-1 Mission Status Report 26, Reference Period 26 Sept – 6 Oct 2014", https://sentinel.esa.int/documents/247904/1484135/Sentinel-1-Mission_Status_Report_26-Period_26_September-6_October_2014

E2015. European Space Agency, Sentinel-1 Mission, https://sentinel.esa.int/web/sentinel/missions/sentinel-1, accessed 30 Nov 2015

F2014. V. Fernandez Arguedas, Q. Zhang and E. Izquierdo. "Multimodal fusion in surveillance applications", *Fusion in Computer Vision*, pp.161-184, Springer

F2015. V. Fernandez Arguedas, "Texture-based vessels classifier for electro-optical satellite imagery", IEEE International conference on Image Processing

G2004. H. Greidanus et al., "Benchmarking operational SAR ship detection", in *Proc. IGARSS 2004*, Anchorage, Alaska, 20-24 Sept 2004

G2006. Harm Greidanus, Naouma Kourti, "Findings of the DECLIMS project – Detection and classification of marine traffic from space", in *Proc. SEASAR 2006: Advances in SAR oceanography from ENVISAT and ERS*, ESA-ESRIN, Italy, 23-26 Jan 2006

G2015. H. Greidanus et al., "Maritime Awareness Systems Performance in the Western Indian Ocean 2014-2015", JRC Technical Report, JRC97935

GS2014. Harm Greidanus, Carlos Santamaria, "First Analyses of Sentinel-1 Images for Maritime Surveillance", JRC Science and Policy Report, JRC92666, EUR 27031 EN, ISBN 978-92-79-44715-0 (PDF), ISSN 1831-9424 (online), doi: 10.2788/132810

K2001. N. Kourti, I. Shepherd, G. Schwartz, P. Pavlakis, "Integrating Spaceborne SAR Imagery into Operational Systems for Fisheries Monitoring", Canadian Journal of Remote Sensing 27(4), pp. 291-305

O1994. T. Ojala, M. Pietikäinen, and D. Harwood, "Performance evaluation of texture measures with classification based on Kullback discrimination of distributions", International Conference on Pattern Recognition, Vol. 1, pp. 582-585

O1996. T. Ojala, M. Pietikäinen, and D. Harwood, "A Comparative Study of Texture Measures with Classification Based on Feature Distributions", Pattern Recognition, Vol. 29, pp. 51-59

O2002. K. Ouchi et al., "Nonuniform Azimuth Image Shift Observed in the Radarsat Images of Ships in Motion", IEEE Transactions on Geoscience and Remote Sensing, Vol. 40, No. 10, pp. 2188-2195

OSM2015. Open Street Map, © OpenStreetMap contributors, https://www.openstreetmap.org/, accessed 27 Nov 2015

PMAR2015. JRC, Web page of PMAR projects, https://ec.europa.eu/jrc/en/research-topic/piracy-maritime-awareness-and-risks, accessed 26 Nov 2015

S2003. J. Sivic and A. Zisserman, "Video Google: A text retrieval approach to object matching in videos", IEEE International Conference on Computer Vision

S2014. M. Stasolla, "Morphological Estimation For Ship SAR Signatures", JRC Scientific and Policy report, JRC91735

SG2015. C. Santamaria and H. Greidanus. "Ambiguity Discrimination for Ship Detection using Sentinel-1 Repeat Acquisition Operations", IEEE International Geoscience and Remote Sensing Symposium (IGARSS) 2015, pp. 2477-2480

SHub2015. ESA, Sentinels Scientific Data Hub, https://scihub.esa.int/, accessed 26 Nov 2015

W1996. P. Wessel and W.H.F. Smith, "A Global Self-consistent, Hierarchical, High-resolution Shoreline Database", J. Geophys. Res., 101, 8741-8743

W2015. https://www.soest.hawaii.edu/pwessel/gshhg/, accessed 27 Nov 2015

Acronyms

AIS	Automatic Identification System
DECLIMS	FP5-funded research project for the detection and classification of maritime traffic from space
EMSA	European Maritime Safety Agency (of the EU)
ESA	European Space Agency
EU	European Union
EW	Extra Wide (a Sentinel-1 imaging mode)
GRD	Ground Range Detected (a SAR image product type)
GRDH	Ground Range Detected - High resolution
GRDM	Ground Range Detected - Medium resolution
GSHHS	Global Self-consistent, Hierarchical, High-resolution Shoreline Database
GT	Ground Truth
HOG	Histogram of Oriented Gradients
IDL	Interactive Data Language (a programming language)
IPF	Instrument Processing Facility
IW	Interferometric Wide (a Sentinel-1 imaging mode)
JRC	Joint Research Centre (a DG of the European Commission)
LBP	Local Binary Patterns
LRIT	Long Range Identification and Tracking
MASE	Programme to Promote Regional Maritime Security
MATLAB	Matrix Laboratory (a programming language)
MMSI	Maritime Mobile Service Identity
MSSIS	Maritime Safety and Security Information System
NCA	Norwegian Coastal Administration
OSM	OpenStreetMap
PMAR	Piracy, Maritime Awareness and Risks
R&D	Research and Development
RCS	Radar Cross Section
SAR	Synthetic Aperture Radar
SLC	Single Look Complex (a SAR image product type)
SM	StripMap (a Sentinel-1 imaging mode)
SUMO	Search for Unidentified Maritime Objects (JRC's automatic ship detection software)
TOPS	Terrain Observation by Progressive Scans
UTC	Coordinated Universal Time
UTM	Universal Transverse Mercator
VDS	Vessel Detection System
VMS	Vessel Monitoring System

Acknowledgments

Sentinel-1 is operated by ESA (European Space Agency) on behalf on the European Commission.

A coastline data base of OpenStreetMap, © OpenStreetMap contributors, was used [OSM2015]. Also the GSHHS coastline was used [W1996, W2015].

MSSIS data were used courtesy of the Volpe Center of the U.S. Department of Transportation, the U.S. Navy and SPAWAR.

Satellite AIS data from NORAIS, AISSat-1 and AISSat-2 were used courtesy of the Norwegian Coastal Administration and the Norwegian Defence Research Establishment (FFI).

LRIT data were used courtesy of the National Competent Authorities of the Flags that participate in the EU LRIT Data Centre, with the help of EMSA (European Maritime Safety Agency).

Satellite AIS data were purchased from exactEarth, Orbcomm / LuxSpace, and SpaceQuest.

European Commission
EUR 27591 EN – Joint Research Centre – Institute for the Protection and Security of the Citizen

Title: Sentinel-1 Maritime Surveillance – Testing and Experiences with Long-term Monitoring

Author(s): Carlos Santamaria, Mattia Stasolla, Virginia Fernandez Arguedas, Pietro Argentieri,
Marlene Alvarez, Harm Greidanus

Luxembourg: Publications Office of the European Union

2015 – 76 pp. – 21.0 x 29.7 cm

EUR – Scientific and Technical Research series – ISSN 1831-9424 (online)

ISBN 978-92-79-53960-2 (PDF)

doi:10.2788/090400

LB-NA-27591-EN-N

JRC Mission

As the Commission's in-house science service, the Joint Research Centre's mission is to provide EU policies with independent, evidence-based scientific and technical support throughout the whole policy cycle.

Working in close cooperation with policy Directorates-General, the JRC addresses key societal challenges while stimulating innovation through developing new methods, tools and standards, and sharing its know-how with the Member States, the scientific community and international partners.

Serving society
Stimulating innovation
Supporting legislation

doi:10.2788/090400

ISBN 978-92-79-53960-2

Publications Office