




JRC TECHNICAL REPORTS

The SUMO ship detection software for satellite radar images

*Short installation and
user guide*

H. Greidanus, P. Argentieri, M. Alvarez,
C. Santamaria, N. Kourti

2017



This publication is a Technical report by the Joint Research Centre (JRC), the European Commission's science and knowledge service. It aims to provide evidence-based scientific support to the European policymaking process. The scientific output expressed does not imply a policy position of the European Commission. Neither the European Commission nor any person acting on behalf of the Commission is responsible for the use that might be made of this publication.

JRC Science Hub

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

JRC107063

PDF ISBN 978-92-79-71881-6 doi:10.2760/349278

Luxembourg: Publications Office of the European Union, 2017

© European Union, 2017

Reuse is authorised provided the source is acknowledged. The reuse policy of European Commission documents is regulated by Decision 2011/833/EU (OJ L 330, 14.12.2011, p. 39).

For any use or reproduction of photos or other material that is not under the EU copyright, permission must be sought directly from the copyright holders.

How to cite this report: H. Greidanus, P. Argentieri, M. Alvarez, C. Santamaria, N. Kourti, *The SUMO ship detection software for satellite radar images – Short installation and user guide*, KJ-01-17-945-EN-N, Publications Office of the European Union, Luxembourg, 2017, ISBN 978-92-79-71881-6, doi:10.2760/349278, JRC107063

All images © European Union 2017; the satellite image background in the cover image © Copernicus 2017.

Contents

Acknowledgements	2
Abstract	3
1 Introduction	4
2 Building the executable	5
2.1 Source code components and license	5
2.2 Building	5
3 Obtaining the executable and running the software	6
3.1 Obtaining and installing	6
3.2 Starting	6
4 Input and output	7
4.1 Input	7
4.1.1 Image input data	7
4.1.2 Auxiliary input data	7
4.1.3 Processing parameters	8
4.2 Output	9
4.2.1 xml output	9
4.2.2 Other output	13
5 User interface	14
5.1 Parameters	14
5.2 Image selection	14
5.3 Visualisation commands	14
5.4 Land masking	15
5.5 Sea ice masking	15
5.6 Ship detection	15
5.7 Saving the detection results	16
5.8 Manual inspection and correction	16
5.9 Next image	18
6 Command line mode	19
6.1 Processing parameters on the command line	19
6.2 Using a configuration file	19
7 Known bugs and shortfalls	22
8 Maintenance	24
References	25
List of abbreviations and definitions	26
Annex. Q&A	27

Acknowledgements

SUMO has been developed over a period of 15 years, and many people besides the authors of this report have worked on it: François-Xavier Thoorens, Iain Shepherd, Guillermo Schwartz, Guido Lemoine, Tony Bauna, Jolyon Chesworth, Elena Aresu, Cristian Cesena, Marte Indregard, Thomas Lefort, Andrea Gabban, Juan Cicuendez. SUMO is also indebted to the many authorities of EU Member States who have been involved in test and validation campaigns over the years.

Authors

Harm Greidanus, Pietro Argentieri, Marlene Alvarez, Carlos Santamaria, Naouma Kourti

Abstract

This is a brief manual for the SUMO software.

SUMO is a software package for semi- or fully-automatic ship detection in satellite SAR (Synthetic Aperture Radar) images, for experimental use. It can work on images from most of the recent and contemporary satellite SARs, across a range of radar frequencies and resolutions. It has been developed at European Commission's Joint Research Centre (JRC) over the last 15 years, and has been extensively tested, demonstrated and validated.

Ship detection from satellite images is used for maritime spatial planning, fisheries control, pollution control, maritime border control, counter-piracy, and other applications in the field of maritime safety, security and sustainability.

The SUMO algorithm, that is behind the software, has recently been described in detail in a journal publication.

This document describes the purpose of SUMO, how to compile the run-time code from the Java source code, how to start and run SUMO in interactive mode as well as in batch mode, and SUMO's input and output data. It also mentions known problems and the necessary maintenance. A list of questions and answers is attached in annex for quick reference.

This manual is published together with the release of the SUMO software as open source under the FreeBSD two-clause license. The software was developed for use at JRC for R&D purposes and is released "as is".

1 Introduction

SUMO is a software package to perform ship detection on satellite SAR (Synthetic Aperture Radar) images. It takes as main input one SAR image and its metadata, and produces as main output a list of detected ships with their attributes.

The SUMO software is experimental. It was made for the purpose of research at JRC, not for operational use, and not with use on a wide range of platforms in mind. It has been constantly in development, adding and changing functionalities. The software source code as released is the current version, which is not bug free. Therefore, not everything described in this manual might immediately work on any platform. However, with the source code available, programmers should be able to make the proper adaptations. The code also contains some sections that are currently inactivated; if they are activated, they enable additional functionalities, that are not covered in this manual.

The SUMO software is the implementation of the SUMO algorithm. The algorithm has been described in detail in Greidanus et al., 2017. The implementation is in Java.

SUMO can run fully automatically, but it is better used semi-automatically. In that case, all the automatically found detections can be viewed on top of the SAR image, and the operator can inspect them and discard any if needed.

Both SUMO's source code and a ready-to-run executable are being made available. This manual will first describe, in chapter 2, how to build the executable code from the source code. This chapter can be skipped for users not interested in the source code; chapter 3 will describe how to install and run the executable code. After that, the input and output data and their formats will be specified. The interactive user interface will be described, followed by the batch mode. Some known shortfalls in the algorithmic approach will be summarised, followed by some known bugs in the present software implementation. At the end of the report are references, a list of acronyms, and some potential questions and their answers.

SUMO is made available under the FreeBSD two-clause license (The FreeBSD Project, 2017).

2 Building the executable

2.1 Source code components and license

Sumo is organised as an Apache Maven project with several subprojects:

- Sumo
 - SumoCore
Contains the classes to manage the platform configuration
 - GeoImageViewer
This project manages the GUI and in particular:
 - Image loading
 - Analysis processes
 - Import/export operations
 - GeoImage
Contains the classes to load all types of images
 - GeoImageAnalysis
Contains the algorithms to analyse the images
 - SentinelXMLReader
This project is a library used by GeoImage to read the Sentinel-1 images
 - GeoLocation
This project is a library that implements a geolocation algorithm. It is used to convert coordinates into image pixel locations and back.

All source components can be found on the repository:

<https://github.com/ec-europa/sumo>.

Third party libraries are:

- Geotools: Open Source Java GIS Toolkit
- Nasa WorldWind: Library for geospatial 3D virtual globe visualisation via OpenGL/WebGL
- JTS: Java Topology Suite
- JOGL: Java OpenGL
- GDAL: Library to read image data
- jhdf5: An alternative Java binding for HDF5
- FenGUI: Old projects that map swing control on OpenGL
- Apache commons
- EclipseLink

2.2 Building

Before building SUMO, one needs to:

- Install Apache Maven
- Install GDAL and h5 third libraries as "local repo" from Maven. To do this, one can use the script in the script folder. This folder contains a version of the libraries compiled for Windows and Linux and the scripts to add them in the local Maven repository.

To build the project it is enough run the following mvn command from the root:

```
mvn clean package install -pl GeoImageViewer -am
```

or one can run the mvn_all.sh script.

3 Obtaining the executable and running the software

3.1 Obtaining and installing

Besides the source code, also the executable can be found on the repository <https://github.com/ec-europa/sumo>.

SUMO does not actually need to be installed. It is enough to copy the SUMO folder that contains the executable onto one's system.

On the github website, click the tab that says 'releases' which will take you to the sub-page '/releases'. Under the headers 'Downloads' there are various software versions available for Windows 7 and Linux. For e.g. the latest Windows version, find the label 'latest release' and click the link 'SUMO_Win64_1.3.5.zip' (1.3.5 being the latest version at time of writing) to download the zipped executable (the zipfile size is mentioned on the right). Unzip the file and copy the resulting directory structure under your 'Program Files' folder.

The Java Runtime Environment (JRE) (or Java Development Kit, JDK) version 1.7.x or above needs to be already installed on the system.

Depending on your system and the presence of dll's, some library paths may need to be adapted. See also section 7, table 6, first entry.

3.2 Starting

SUMO is started by running the 'start.bat' command file, probably easiest from a shortcut copied (by the user) onto the desktop.

SUMO comes up with two windows, a command window for process tracking and error display, and the main display window. The main window can be sized as desired by dragging the edges. It has two tabs on top left, 'Image' and '3D'.

The first time SUMO is started, it may be necessary for proper initialisation to select the '3D' tab, and then to resize the window (by dragging a corner), and drag the globe that should be visible around with the mouse. This should result in having the globe in a centred position. Then the 'Image' tab should be selected, and after that SUMO should be closed by selecting 'System → Exit' or simply by closing either the main window or the command window with the usual close button (x) on top right. After that, SUMO can be started again and should then be properly initialised.

4 Input and output

4.1 Input

4.1.1 Image input data

This concerns the satellite SAR image (plus its metadata) that should be analysed. SUMO accepts images from the satellites and sensors listed in the table below.

Table 1. Satellites / sensors accepted by SUMO.

Satellite / Sensor	2-letter code	Format	Remark	File to select
RADARSAT-1	R1	CEOS		DAT_01.001 or equivalent
RADARSAT-2	R2	tiff+xml		product.xml
TerraSAR-X (and TanDEM-X)	TX	tiff+xml	SLC data are in .cos format which is presently not readable	*.xml
Cosmo-SkyMed	CK	HDF5		*.h5
ALOS2-PALSAR	A2	CEOS; tiff	Only CEOS gives full functionality	summary.txt
ERS-1,-2	E1, E2	CEOS		DAT_01.001 or equivalent
ENVISAT-ASAR	AS	CEOS		*.N1
Sentinel-1	S1	tiff+xml		manifest.safe

Some of these systems give the image product as a single file, for others one image product consists of a set of files, sometimes in a directory structure. The last column indicates the filename that should be selected (by clicking) in order to open the image (see section 5.2).

The image should be in natural (range–azimuth) coordinates. It should NOT be geocorrected, DEM-corrected, orthorectified or map-projected. It should however be georeferenced, i.e., the image metadata should specify the relation between pixel coordinates and geographical coordinates.

The image should ideally be in ground range, amplitude and multi-looked, which is one of the normal SAR image product types, variously called GRD (Ground Range Detected), MGD (Multi-look Ground range Detected), Level 1, Level 1B or Level 1.5. In addition, it is possible to ingest Single Look Complex (SLC) images, which are in slant range and have the complex amplitude. However, SLC images are upon ingestion directly converted to detected (absolute) amplitude without the upsampling that would be needed to preserve Nyquist sampling, which means their information content is not fully used.

4.1.2 Auxiliary input data

Coastline data

In order to mask out the land that may be inside an image, SUMO uses a global coastline vector file. There are several prepared options from which the user can select. They include GSHHG (Wessel and Smith, 1996 and 2016) and OpenStreetMap (2017). In

addition, coastline files have been prepared by seaward buffering of the OpenStreetMap coastline by 50 m, 100 m and 250 m, and these vector files can also be selected.

The use of a buffered coastline reduces false alarms due to parts of the land that are not accurately covered by the original coastline vector file. It is also faster because the buffering results in smoothing of details and fewer vertices.

All global coastline vector files are in shapefile format, and they have been cut up in tiles to speed up ingestion.

Finally, the user can ingest his own coastline vector file. For the easiest way to do so, the user should put his coastline vector file in shapefile format in a sub-directory under <SUMODir>\resources\coastline (where <SUMODir> is the main directory of the SUMO software), next to the sub-directories where the above mentioned coastline vector files are found. Alternatively, the shapefile can be stored anywhere on disk and loaded into SUMO by selecting it in the user interface. (See section 5.4.)

Sea ice cover

In order to mask out areas that are covered by sea ice (which would give rise to false alarms, like on land), a vector file can be ingested that maps out the sea ice cover extent. The vector file can be available locally, or can be downloaded at runtime from a remote website. At present, the only remote website that SUMO can deal with is that of the U.S. National Ice Center / Naval Ice Center (U.S. NIC, 2017) which hosts a daily updated ice map. (See section 5.5.)

Look-up tables

For fast operation, SUMO uses look-up tables to convert statistics that are measured from the image (mean and standard deviation) to a Constant False Alarm Rate (CFAR) detection threshold, given the Equivalent Number of Looks (ENL) of the image. These tables have been pre-computed by numerical integration of the K distribution probability density function, for a wide set of ENL values. (For details on the algorithm and the concepts CFAR, K distribution, ENL, etc., please see Greidanus et al. 2017.) The look-up tables are pre-supplied and the user does not need to worry about them.

However, it could be that in the future, satellite SAR image products become available with new ENL values, for which no look-up tables were included. In that case, additional look-up tables would need to be supplied. As a workaround, a nearby ENL value for which a look-up table is available can be used, with an adjusted detection threshold. If an ENL value is used in the detection run that is higher than the actual ENL of the image, then the detection threshold adjustment should be raised; and vice versa.

4.1.3 Processing parameters

Although SUMO is designed to run automatically without adjustments on any satellite SAR image, it is nonetheless still necessary to use some processing parameters in order to be able to deal with the widest possible collection of images under the widest possible range of circumstances. The table below lists the parameters that can be set for a detection run. See section 5.6 for how to specify them.

Table 2. Adjustable parameters for a detection run.

Parameter	Normal value	Remarks	More detections (both real and false) when ...
Nominal false alarm rate	10^{-7}	Another value that the 'normal' is usually not needed, as the detection threshold can be effectively adjusted by the next parameter	lowered
Detection threshold adjustment	1.5 for co-pol (HH and VV) 1.2 for cross-pol (HV and VH)	This is the parameter that may need to be changed by the user most frequently, in most cases by raising it from the 'normal' value in order to prevent too many false alarms in cases of heavy sea clutter	lowered
ENL	Automatically read from the image metadata or derived from the image type	A user-specified value overrides the automatically derived value	raised
Land mask seaward buffer size	0 m (if using pre-buffered land mask) 100 m (otherwise)	The global coastlines that are used can be inaccurate in some parts of the world, necessitating a larger buffer. Very complex coastlines are processed faster with a larger buffer.	lowered

4.2 Output

4.2.1 xml output

The ship detection result can be written out in several formats. The main one is xml. How to select this format and write out the results is explained in section 5.7.

A default output filename is proposed by the software with the following structure:

```
SS_yyyymmdd_HHMMSS.xml
```

where SS is a two-letter code for the sensor according to table 1, and the rest is the date and time (UTC) of the image start, obtained from the image metadata. So an instance of the default output filename could be e.g.:

```
S1_20160517_093625.xml
```

The user interface allows to modify this filename before saving.

The first level hierarchical structure of the xml output file has three sections, one on the image, one on the analysis parameters, and one on the detection results, as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<analysis>
  <sat_image_metadata> ...
</sat_image_metadata>
  <vds_analysis> ...
</vds_analysis>
  <vds_target> ...
</vds_target>
</analysis>
```

Image metadata

The section on image information has the following contents (here, a specific example is quoted):

```
<sat_image_metadata>
  <ImId>S1_20160517_093625</ImId>
  <image_name>S1A_S2_GRDH_1SDV_....SAFE</image_name>
  <timestampStart>2016-05-17 09:36:25.004626</timestampStart>
  <sensor>S1</sensor>
  <pol>VV VH</pol>
  <polnumeric>4,3</polnumeric>
  <TimeStart>2016-05-17 09:36:25.004</TimeStart>
  <TimeStop>2016-05-17 09:36:44.051</TimeStop>
  <gcps>
    <gcp>
      <row>1</row>
      <column>1</column>
      <lon>125.0852944840508</lon>
      <lat>24.068215398846647</lat>
    </gcp>
    <gcp>
      <row>1</row>
      <column>8078</column>
      <lon>125.86467264222956</lon>
      <lat>24.210395600412895</lat>
    </gcp>
    <gcp>
      <row>12818</row>
      <column>1</column>
      <lon>124.82884663741073</lon>
      <lat>25.218184848801602</lat>
    </gcp>
    <gcp>
      <row>12818</row>
      <column>8078</column>
      <lon>125.61542146362328</lon>
      <lat>25.36013515420748</lat>
    </gcp>
  </gcps>
</sat_image_metadata>
```

The field <ImId> receives the same value as the default proposed output filename.

The field <image_name> receives either the filename of the image or the directory name of the image, depending which of the two is the more informative (which depends on the sensor).

The field <timestampStart> is copied from the metadata and can have a different format or length depending on the sensor.

The field <sensor> encodes the sensor in two letters as for the default proposed output file name (table 1).

The field <pol> has the list polarimetric bands in the image, from HH, HV, VH, VV, separated by a space.

The field <polnumeric> encodes the polarimetric bands as comma-separated numbers (1=HH, 2=HV, 3=VH, 4=VV).

The fields <TimeStart> and <TimeStop> are taken from the image metadata, but are always converted to yyyy-mm-dd HH:MM:SS.MMM. The values in <timestampStart> and <TimeStart> are the same except for rounding, and should correspond to the UTC of the first azimuth line in the image.

The numbers under <gcps> give the locations of the four image corners, both in row/column and in latitude/longitude. This means that also the number of pixels in the image can be found here.

The nomenclature and display used for the two directions in the image is:

x = columns = horizontal = left-to-right = range

y = rows = vertical = top-to-bottom = azimuth

Analysis parameters

The section on analysis information has the following contents (again, a specific example is quoted):

```
<vds_analysis>
  <run_time>2017-06-16 15:17:10.270</run_time>
  <run_version>Auto12</run_version>
  <run_version_num>1</run_version_num>
  <detector_version>SUMO_1.3.4</detector_version>
  <algorithm>k-dist</algorithm>
  <parameters>27.0,0.0, 0.0, 1.2, 1.5,0.00</parameters>
  <nr_detections>2</nr_detections>
  <buffer>0</buffer>
  <land_mask_read>OSMLandPoly_20141001_250m.shp</land_mask_read>
</vds_analysis>
```

The field <run_time> receives the computer clock time when the xml file was saved.

The fields <run_version> and <run_version_num> receive the values that the user entered in those fields in the user interface when saving the file (section 5.7).

The field <detector_version> receives the name and version number of the software.

The field <algorithm> receives the user choice made for the algorithm (presently only k-dist).

The field <parameters> contains the numerical values of the following analysis parameters, in sequence: ENL, th_HH, th_HV, th_VH, th_VV. The values of these parameters are assigned before starting a detection run (section 5.6).

The field <buffer> contains the seaward buffer applied to the land mask in units of pixels as assigned at the same occasion.

The field <land_mask_read> specifies the filename of the coastline vector file that was used for land masking.

The field <nr_detections> contains the number of targets that were detected.

Detected targets

The section on the detected targets has the following contents (again, a specific example is quoted):

```
<vds_target>
  <boat>
    <target_number>1</target_number>
    <lon>125.620649</lon>
    <lat>24.821851</lat>
    <xpixel>6970.0</xpixel>
    <ypixel>7031.0</ypixel>
    <detecttime>2016-05-17 09:36:25.004</detecttime>
    <inc_ang>28.024</inc_ang>
    <reliability>0</reliability>
    <nr_pixels>133</nr_pixels>
    <length>200.0</length>
    <width>64.0</width>
```

```

    <size_class>L</size_class>
    <heading_range>-73.0</heading_range>
    <heading_north>0.0</heading_north>
    <lwh_reliability>>false</lwh_reliability>
    <manual_add>>false</manual_add>
    <noise>0</noise>
    <significance>0.0,0.0,0.0,33.482</significance>
    <threshold_tile>0.0,0.0,0.0,784.681</threshold_tile>
    <max_value>0,0,0,1354</max_value>
  </boat>
  <boat> ...
</boat>
...
</vds_target>

```

One item `<boat>` is included for each detected target, to the total of the `<nr_detections>` from the analysis section.

The field `<target_number>` has a sequential number from 1 to `<nr_detections>`.

The fields `<lon>` and `<lat>` contain the geographic position of the target centre, in decimal degrees, WGS84.

The fields `<xpixel>` and `<ypixel>` contain the pixel location of the target centre in the image; `x` = range, `y` = azimuth, see Image metadata above.

The field `<detecttime>` has the UTC time that the target was detected as yyyy-mm-dd HH:MM:SS.MMM. This varies between the image start and stop times according to the azimuth position of the target in the image.

The field `<inc_ang>` gives the target's incidence angle in degrees.

The field `<reliability>` currently can have the following two values:

0: likely real target;
3: likely false alarm.

The field `<false_alarm_cause>` is only present for targets that are likely false alarms, e.g. `<false_alarm_cause>AA</false_alarm_cause>` for azimuth ambiguities, or TP for a secondary azimuth ghost that can occur in Sentinel-1 images.

The field `<nr_pixels>` has the number of pixels in the target.

The fields `<length>` and `<width>` have those estimated target attributes in meters.

The field `<size_class>` is S, M or L directly derived from `<length>`.

The fields `<heading_range>` and `<heading_north>` are the estimated target heading in degrees with respect to the range direction and with respect to North, respectively.

The field `<lwh_reliability>` is true if the parameters length, width and heading are deemed to be reliable (e.g. for a large, well-defined, elongated target) and false if not.

The field `<manual_add>` is true if the target was manually added (see section 5.8).

The field `<threshold_tile>` is the detection threshold of the tile in which the target is located, as a row of four numbers corresponding to polarisation bands HH, HV, VH, VV in that order; bands that do not occur have 0.

The field `<max_value>` is the maximum pixel value that occurs on the target, for the polarisation bands HH, HV, VH, VV respectively (bands that do not occur have 0).

The field `<significance>` is the number of standard deviations above the background mean that the target maximum reaches, i.e.,

$\text{<max_value> = <bgnd_mean> + <significance> * <bgnd_stdev>}$,
as a row of four numbers corresponding to polarisation bands HH, HV, VH, VV (0 for bands that do not occur).

(NB: <bgnd_mean>, the mean of the background, and <bgnd_stdev>, its standard deviation, are not output.)

The field <noise> is currently not used.

4.2.2 Other output

Besides in xml format, the detection results can also be exported as comma-separated (.csv), shapefile (.shp + adjoined files), geography markup language (.gml) or Google Earth (.kmz) file; however, the output in these formats is not complete and / or not fully tested for correctness.

5 User interface

5.1 Parameters

After SUMO is started following section 3.2, it is possible to change a series of default parameters under the drop-down menu 'System → Preferences'. The drop-down menu is on the left in the top bar of the user interface. However, these can be changed any time and the default values are good to start with.

5.2 Image selection

The first step of analysing an image is to select it. This is done by clicking 'Import → Image' from the drop-down menu. The choice for 'data type' should be left to 'image'. Pressing 'OK' opens a window where the user can select the image to ingest by clicking on the file according to the rightmost column of table 1. SUMO will read the image metadata, open the image file and display the top left part of the image on the screen at 1:1 scale (showing each pixel). The status bar on the bottom right will indicate the ingestion activity. Under the menu 'Layers' on bottom left, the name of the image file will appear.

5.3 Visualisation commands

At the right of the Layers menu on bottom left, there is a box for visualisation commands. Clicking the mouse in this box opens it for typing. The possible commands are listed in table 3.

Table 3. Visualisation commands for the command box.

Command	Meaning	Function
h	home	Show entire image.
bs	band switch	Switch between bands in a dual-pol image. In case of more than 2 bands, cycle through the bands.
c +50	contrast	Increase brightness by 50 units (use any value). Use – to decrease the brightness. For the + sign, use the numerical keypad of the keyboard.
empty (i.e., hit only enter)		Repeat previous command.

So, to display the entire image, type 'h' in the command box (followed by enter). The image will take a while to load, tile by tile. Again, the status bar on the bottom right will show the progress of the loading.

If the image is dual polarisation, the other band can be loaded for viewing by typing 'bs' in the command box. (If this is typed before the previous band is fully loaded, it will have no effect.) If the image has more than two bands, repeat the 'bs' command to load the other bands for viewing.

By only giving an enter on the keyboard, the last command is repeated. Therefore, after having given the command 'bs' once, subsequent empty commands (just enter) will switch between bands (if dual pol) or cycle through bands (if more than two bands). Likewise, after e.g. a 'c +20' command is given, it can be repeated as often as needed to slowly increase the brightness by hitting enter repeatedly.

The image can be zoomed in and out by using the roller on the mouse. It can be panned by pressing the shift key on the keyboard while moving the mouse around.

SUMO creates small png image files in a cache on disk of those parts of the original image that are being displayed. Different pngs are created for different sub-areas and different zoom levels. It takes a while to make a png, but when a previously displayed area and zoom level are re-displayed, the cached png file is used, so then it is much faster. This is why the first sequence of 'h', 'bs' takes some time, but then subsequent enter commands (repeating the 'bs') switch instantaneously between the bands.

Overview box

In the top left, on top of the image, is a small box that contains an overview of the entire image. The currently zoomed-in part of the image is indicated there with a green rectangle. Clicking the cursor somewhere in the overview box immediately centres the zoom on the click location, obviating the need for panning there.

Right-clicking on the overview box makes it disappear.

5.4 Land masking

If there is land in the image, it should be masked out before running the ship detection. To load a coastline for land masking, go under 'Import → Coastline' in the drop-down menu and select one of the provided vector coastlines. The OpenStreetMap coastline is provided in pre-buffered versions that are called 'OSMLandPoly_date_B' with B the seaward buffer size in meter. The vector coastline should load quickly, unless the coastline is very complex in which case it takes a bit longer.

A coastline can also be loaded as a shapefile from any location on disk using 'Import → Vector' with Type = shp and Mask type = coastline.

It is possible to load more than one vector coastline at this stage. Each vector coastline shows up in the display and appears as a line under the bar 'Layers' in the bottom left.

Layers (I)

All the lines under the 'Layers' bar on bottom left can be left-clicked to activate or deactivate. Their content will accordingly be displayed or not. Clicking the line that identifies the image itself (the top line at this point) deactivates the entire image and all its associated layers. Deactivating in this way does not remove the layer(s) – for that, see below under Layers (II).

5.5 Sea ice masking

The shapefile for ice cover can be imported from a local file or downloaded from the remote U.S. National Ice Center (NIC) service web site. Use 'Import → Vector' with Type = shp and Mask type = ice.

To use the NIC ice mask, tick the box 'Load from remote site'. The NIC repository contains one shapefile for each day, whose name contains the date of the day. The URL to download the correct shapefile for the analysis is built using the image date.

Leaving the 'Load from remote site' box unchecked will open a file explorer dialogue to choose a local shapefile.

5.6 Ship detection

To perform the ship detection on the image, use 'Analysis → VDS' from the drop-down menu. (Note that the 'h', 'bs' and other commands discussed before are only for visualising the image. It is not necessary to do any visualising before starting the ship detection process.) A pop-up menu appears. On top, the vector coastline can be selected from the ones that have been loaded. Below that, an ice cover mask can be selected. The selection on the next line is for algorithm; however, only the K distribution algorithm is currently implemented. The following lines are for the detection threshold adjustment, one line per band in the image. The normal values to use are 1.5 for co-pol bands (HH or VV) and 1.2 for cross-pol bands (HV or VH). These should be raised if the clutter level is

very high (e.g., many waves or other inhomogeneities in the image) or if the interest is only in strong targets (big ships). They could e.g. be raised by 0.3, i.e. to 1.8 (co-pol) / 1.5 (cross-pol), to prevent the detection of very small targets and weak false alarms. On the next line, a seaward buffering distance for the land mask can be chosen in units of pixels. If a pre-buffered land mask was already selected, this can normally be left at 0. If no pre-buffered land mask was selected, it is useful to introduce the buffer here. It is however faster to use the pre-buffered land mask. It is possible to combine a pre-buffered land mask (to a distance in meters) with an additional buffering here (to an additional distance in pixels). The final line is to specify the Equivalent Number of Looks (ENL) of the image. The default value is in most cases derived from the input image so should be correct. In some cases, however, it may be wrong, and then the correct value should be entered.

Clicking 'OK' at the bottom of this interface starts the ship detection process. Its progress is indicated on the bottom right progress bar. It can also be followed in the command window. Each polarimetric band of the image is analysed in turn. For each band, first the image is scanned for CFAR detected pixels. Then all detected pixels are clustered into objects. After these two steps have been done for each band separately, in the case of a multi-band image, there is an additional step that merges all detected pixels from all bands and re-does the clustering, in that way merging the ship detection results from all bands into one single detection result.

Layers (II)

All detection results show up as a line in the 'Layers' column. So, e.g., for a two-band HH+HV image, there will be three lines: 'VDS analysis HH', 'VDS analysis HV', and 'VDS analysis all bands merged'. The single band result names are followed by the value of the detection threshold adjustment that was used. By left-clicking on such a line, it can be activated (displayed) or deactivated (not displayed). The detections in each band are displayed as a symbol on the image. The size, shape and colour of the symbols can be changed under the 'System → Preferences' menu.

By right-clicking a line in the 'Layers' column, a dialogue box concerning that layer pops up. This dialogue box contains tabs that enable a number of actions on the layer, such as changing the display style, editing or saving. The first tab provides information on the layer, and has also a 'Delete' button to remove the layer. Removing the layer does not delete any data on the computer disk; it means that the layer disappears from the display, and will have to be re-ingested or re-computed if needed again. Close the dialogue box with the close button on its top right.

5.7 Saving the detection results

Right-clicking one of the layers 'VDS analysis' is needed to edit and save the ship detection result. For multi-band images, it is usually only needed to edit and save the 'VDS analysis all bands merged' layer, and not the results of the individual polarisation bands. After right-clicking the VDS analysis layer to save, click the 'Save' tab in the pop-up box to, at this point, save the fully automatic ship detection result before any manual editing. Select which 'Type' of output file to save; the normal type is 'sumoXML'. Click the 'Choose' button to specify the output directory and file name of the ship detection results. A default filename is provided which can be modified or replaced. Optionally, a string can be entered in the field 'Run version' and a number in the field 'Run version Number'; these attributes will be copied into the output xml file. One can write e.g. 'Auto' in the 'Run version' field to indicate that this is the automatic, unedited result. Click 'Save' to write out the ship detection results file.

5.8 Manual inspection and correction

Clicking the 'Data' tab of the pop-up box shows a list of all detected targets, one per line, with their attributes. By clicking on one of the lines, that target is selected, and the display jumps there to show the target and its surroundings. At this point, SUMO is in

manual target inspection mode and responds differently than before. It reacts to the pressing of individual keys on the keyboard, as specified in table 4 below.

Table 4. Key functions in manual target inspection mode.

Key	Function
z	Zoom in (can be repeated)
Z	Zoom out (can be repeated)
b	Switch or cycle between bands
d	Toggle display of the extracted target outline and a big box around it, resetting zoom level
a	Toggle display of range-azimuth cross
e	Edit target attribute values
arrow down	Select next target in the list
arrow up	Select previous target in the list
backspace	Delete target
control-z	Undo last delete (in reverse sequence)

A thorough visual inspection of the results of an automatic detection run can be done by selecting the first target, and after inspecting it press the down arrow to go to the next target, etc. At each target, the z / Z keys can be used to zoom in / out, b to switch bands, and backspace to delete a target that is deemed a false alarm. After deleting a target, the next target is automatically selected so the arrow down key should not be used. If a target is accidentally deleted, it can be recovered by pressing control-z. This also works to recover more than one deleted target. However, they are re-inserted in the target list at the end.

In the pop-up box, below the target list, a number of lines appear with checkboxes. Ticking these boxes shows, respectively, the tiles that were used in the CFAR detection process, the buffered land mask that was used, suspected azimuth ambiguities, and several selections of pixels that were used in the detection process. Regarding the suspected azimuth ambiguities, ticking the line 'azimuthambiguities' will highlight them in red, and these are the same that will have a value of 3 in the field <reliability> in the output xml file (section 4.2.1, Detected targets). Note that any of these ticked items will only be actually visible if the layer is active, as toggled by left-clicking the layer as described above under section 5.6, Layers (II).

In the manual target inspection mode, it is not possible to do the visualisation commands and actions that have been described before, such as activating / deactivating layers, zooming / panning with the mouse, or giving commands in the command box. In order to do that, the manual target inspection mode should be left by closing the pop-up box (button top right). The manual target inspection mode can be re-entered by again right-clicking the desired 'VDS analysis' layer, selecting the 'Data' tab and selecting a target line.

The 'Edit' tab in the pop-up box of a selected analysis layer has a checkbox 'Interact' that can be clicked to interactively edit targets, if needed. If this box is checked, a new pop-up window appears with the three buttons 'Add', 'Del' and 'Move'. This new box may be behind the already displayed pop-up box, which anyway needs to be closed to access the

'Add-Del-Move' buttons. Now we have left the manual target inspection mode. To manually add a target, click the 'Add' button and move the cursor (by panning / zooming and pointing with the mouse) to the location in the image where the new target needs to be added. Left-click the mouse on that location; a new dialogue box opens where attributes of the new target may optionally be entered. The new target is actually added after pressing 'OK'. This can be repeated. Re-click the 'Add' button to exit the adding mode. To delete a target (in an alternative way than previously described), click the 'Del' button, and move in the image (by panning / zooming and pointing with the mouse) to the location of an existing target to be deleted. Clicking the centre of the target once with the cursor will highlight it. Double-clicking will delete it. After finishing these operations, the small 'Add-Del-Move' pop-up box cannot be closed with its close button; instead, the layer has to re-opened by right-clicking it (the 'Edit' tab should still be the open one) and the 'Interact' checkbox has to be unchecked.

Note that any layer may be visible or not depending on whether it is active; *independently* of that, one layer may be selected by right clicking it, enabling the edit / target select / save options. So, the layer that is being edited or in which a target is selected is not necessarily one of the active (displayed) layers.

To save the ship detection results after manual editing, the same action is needed as described above. Under the 'Save' tab in the pop-up box from the right-click-selected layer, choose a file name – different from the file name that was used to save the fully automatic result, if that was saved –, and possibly type e.g. 'Manual' in the 'Run version' field. The pop-up box can then be closed with its close button.

This finalises the analysis of the image.

5.9 Next image

Before analysing the next image, the layer with the current image name (top layer) can be left-clicked to deactivate it. However, in that case the layers remain in the layers list on the screen. To fully close the image and clean the 'Layers' stack, right-click the layer with the image name to bring up its pop-up box, and click 'Delete' (under the 'Description' tab). (Obviously this does not delete the image on the disk.)

The next image can be selected as before.

6 Command line mode

SUMO can also be run in command line mode. There are two ways to do it: passing the processing parameters directly on the command line or creating a configuration file.

6.1 Processing parameters on the command line

This method creates an instance of SUMO that analyses one single image. It can be used on small scale experiments (e.g. analysing a few images, one at a time, on a desktop PC), but it is primarily intended for large scale campaigns where multiple such instances of SUMO run on a multi-node processing platform.

This method does not accept parameters for sea ice masking. One way around this issue is to combine, outside SUMO, the coastline shapefile and the required sea ice mask, and to input the combined shapefile as coastline shapefile.

The syntax is as follows:

```
<sumo_batch_script> -i <input_image> -thh <threshold_hh> -thv  
<threshold_hv> -tvh <threshold_vh> -tvv <threshold_vv> -sh  
<coastline_shapefile> -b <coastline_buffer> -o <xml_output_folder>
```

where:

<sumo_batch_script> is 'start_batch.bat' for Microsoft Windows, 'start_batch.sh' for Linux.

<input_image>: filename (with folder) of the input image to be analysed by SUMO (last column in table 1)

<threshold_xx>: floating point number with the detection threshold adjustment to be used in the ship detection analysis (xx can be hh, hv, vh or vv).

<coastline_shapefile>: filename (with folder) of the shapefile (.shp) to be used for land masking.

<coastline_buffer>: integer with the buffer (in pixels) to be applied to the coastline.

<xml_output_folder>: folder name where the SUMO output (xml files) will be stored

Example:

```
start_batch.bat -i C:\somefolder\S1A_IW_GRDH_1SDV_20151122T053659_  
20151122T053724_008713_00C679_C5E8.SAFE\manifest.safe -thh 1.5 -thv 1.2  
-tvh 1.2 -tvv 1.5 -sh C:\somefolder\Sumo\Sumo_1.3.5\  
resources\coastline\OSMLandPoly_20141001_250\OSMLandPoly_20141001_250m.s  
hp -b 0 -o C:\somefolder\output\test-20170627\
```

Note that this example uses a dual-pol VV+VH product. The parameters thh and thv will therefore be ignored.

After running this command, the output folder will contain an xml file (S1_20150626_052736.xml) with the SUMO results for the two polarimetric bands merged. A subfolder (S1A_IW_GRDH_1SDV_20150626T052736_20150626T052801_006540_008B28_EFC0.SAFE) will contain the results for each one of the bands (S1_20150626_052736_VV.xml and S1_20150626_052736_VH.xml) and, again, for the bands merged (S1_20150626_052736_Merged.xml).

6.2 Using a configuration file

This method passes all the parameters to SUMO via a configuration text-file. It is intended for midsized campaigns run on a desktop PC, since a single instance of SUMO will be created and this instance will analyse sequentially the images specified in the configuration file. This method handles sea ice masking. To use sea ice cover shapefiles, parameters were added in the configuration file to specify the repository and the

shapefile for the analysis. When using a remote repository, the ice shapefile is downloaded before analysing the image.

The syntax is as follows:

```
<sumo_batch_script> -gconf <config_file>
```

where:

<sumo_batch_script> is 'start_batch.bat' for Microsoft Windows, 'start_batch.sh' for Linux.

<config_file>: a text-file configuration file that accepts the following parameters.

Parameters (each parameter in a new line; empty lines and lines starting with the # symbol are ignored):

- txx: floating point number with the detection threshold adjustment (xx can be hh, hv, vh or vv)
- shape_file: filename (with folder) of the shapefile (.shp) to be used for land masking
- buffer: integer with the buffer size (in pixels) to be applied to the coastline
- use_file_list: it indicates whether the images to be analysed are provided as a list in a text-file. Possible values:
 - o true: the list is in the 'input_file' parameter
 - o false: the images will not be taken from a list in a text-file but from a directory folder (see 'input_folder' parameter)
- input_file: name of the text-file with the list of images to be analysed. In the text-file, one image (full path) per line. To be used in combination with 'use_file_list'
- input_folder: SUMO will analyse all the images in this folder, unless indicated otherwise by the 'filter' parameter
- filter: wildcard character pattern. Optional parameter. To be used alongside the 'input_folder' parameter. Only the images whose folder names match the pattern will be analysed by SUMO. Example pattern: CSKS1*
- output_folder: folder where the output xml files will be saved. Each image analysed will have a subfolder and the xml files for that image (one per polarimetric band and the merged one) will be stored in that subfolder
- xml_output_folder: in partial duplication to the above, a folder where only the xml files with the detection results of the polarimetric bands merged will be saved
- run_version: string that will be saved in the <run_version> field of the output xml files (see section 4.2.1)
- run_version_number: integer that will be saved in the <run_version_num> field of the output xml files (see section 4.2.1)
- use_ice – Possible values:
 - o none: analysis without ice mask
 - o single: use only one shapefile for all images
 - o repository: use a remote or local repository to find the correct ice mask shapefile based on the image date
- ice_repository_site – Possible values:
 - o local: use a local repository
 - o remote: use a remote repository
- ice_repository_path – Possible values:

- local: use a local repository
 - *for a single shapefile*: an absolute path
 - *for a remote or local repository*: a url that can contain a Java date pattern between “%” which will be replaced with the correct date. Example: http://www.natice.noaa.gov/pub/daily/arctic/%yyyy%/ice_edge/
- ice_pattern_name – Possible values:
- a filename that can contain a Java data pattern like %yyyyDDD% where yyyy=year and DDD=day of year. Example: nic_autoc%yyyyDDD%n_pl_a.zip

Example of a configuration file:

```
#Comment
thh=1.5
thv=1.2
tvh=1.2
tvv=1.5

shape_file=C:\\somefolder\\OSMLandPoly_20141001_250m.shp
buffer=0

use_file_list=false
input_folder=C:\\somefolder\\images
output_folder=C:\\somefolder\\output
xml_output_folder=C:\\somefolder\\output\\xml
filter=*

run_version=BATCH
run_version_number=1

#Default is none
#none=analysis without ice
#single=use only one shp file for all images
#repository=use a remote or local repository to find the correct
# ice shp file based on the image date
use_ice= repository
ice_repository_site=remote
ice_repository_path=http://www.natice.noaa.gov/pub/daily/arctic/%yyyy%/ice_edge/
ice_pattern_name=nic_autoc%yyyyDDD%n_pl_a.zip
```

7 Known bugs and shortfalls

A number of issues with automatic ship detection in SAR images remain unresolved in most existing ship detectors, including in SUMO. Some were discussed in H. Greidanus et al., 2017. The table below lists the main ones.

Table 5. Known shortfalls in the SUMO detection algorithm.

Shortfall	Impact	Remedy
Not all land (including islands and reefs) is masked out because the coastline is not accurate	False alarms	Use larger buffer with coastline; or delete false alarms by hand
Ship wake is detected as a target, especially in cross-pol	False alarm	Delete by hand
Areas of extremely low (zero) backscatter	Many false alarms at the edges of these areas	Delete false alarms by hand; or consider the image in question not suitable for ship detection
Situations of very high wind and waves	Many false alarms from breaking waves	Raise detection threshold; or delete false alarms by hand
Spurious detections, on radio interference, local ocean / atmosphere effects, range or azimuth ambiguities from a source outside the image, icebergs	False alarms	Delete false alarms by hand inasmuch as recognised as such
Only one value of the ENL can be set for an image, although some ScanSAR scenes have a different ENL for each sub-swath	Too many false alarms, or missed detections, in one or more sub-swaths of a ScanSAR scene	As the ScanSAR image is a product that is mosaicked from sub-swaths, it is not possible to use a correct ENL per sub-swath. A partial remedy is to change the ENL value in the user interface, to get correct results but only for some sub-swaths at the time

The SUMO software has been constantly developing and changing over the years. The version that is released with this manual has a number of known bugs, which are listed in the table below. As opposed to the shortfalls which are imperfections of the detection algorithm, the bugs are imperfections of the software implementation. These bugs do not always occur but may occur in some cases.

Table 6. Known bugs in the software implementation.

Bug	Impact	Remedy
GDAL libraries not found (a warning in the command window when starting up)	Cannot properly open various image types	Install necessary dll's on your system. Adapt search paths in start.bat file. Possibly, install Microsoft Visual C++ 2010.
Viewing tiles are wrongly displayed	In the viewer, some tiles are shown with the wrong resolution or not at all	Pan the viewer to correctly visualise the desired area; or empty the cache (chapter 8); or increase the 'Max Number of Tiles' under 'System → Preferences'
Missed targets	Some targets are not detected even though they are strong	Manually add the missed targets
Error in land masking	Some false alarms are detected on land and/or some true targets at sea are missed	Use other coastline
Under 'Analysis → VDS', the default values for co-pol and cross-pol thresholds are interchanged	If left unchanged, would lead to wrong detection thresholds	Change manually on user interface
Under 'Analysis → VDS', the default value for ENL is set to 1.0 because the correct one was not found in the metadata	Detection threshold is too high, leading to missed targets	Change manually on user interface
One strong target is detected multiple times (mostly twice)	False alarm	Delete one detection by hand

8 Maintenance

As mentioned, SUMO creates small png image files in a cache on disk of those parts of the original image that are being displayed. The name of the cache directory is specified under 'System → Preferences' as 'Cache Folder' and can be changed. All created png images are kept, so that when a previously viewed location and zoom level of any image is re-visited, it is very quickly visualised. However, this also means that the cache folder keeps growing. If it takes too much disk space, it can be emptied by 'Tools → ClearCache', or outside SUMO in the file manager.

SUMO creates a directory with persistent information under the user directory with the directory name '.sumo'. (The user directory is the one above the 'My Documents' directory in Windows.) If there are problems with getting SUMO to run, it could help to delete this .sumo directory (which will be automatically re-generated).

References

The FreeBSD Project, <https://www.freebsd.org/copyright/freebsd-license.html> (accessed 12 June 2017)

H. Greidanus, M. Alvarez, C. Santamaria, F.-X. Thoorens, N. Kourti, P. Argentieri, 'The SUMO Ship Detector Algorithm for Satellite Radar Images', *Remote Sensing* **2017**, 9(3), 246; doi:10.3390/rs9030246

OpenStreetMap, <https://www.openstreetmap.org/> (accessed 12 June 2017)

P. Wessel, W.H.F. Smith, 'A Global Self-Consistent, Hierarchical, High-Resolution Shoreline Database', *J. Geophys. Res.* **1996**, 101, 8741-8743

P. Wessel, W.H.F. Smith, 'A Global Self-consistent, Hierarchical, High-resolution Geography Database', <http://www.soest.hawaii.edu/wessel/gshhg/> (accessed 12 June 2017)

U.S. National Ice Center, <http://www.natice.noaa.gov/> (accessed 12 June 2017)

List of abbreviations and definitions

CEOS	Committee on Earth Observation Satellites
CFAR	Constant False Alarm Rate
ENL	Equivalent Number of Looks
GRD	Ground Range Detected
GUI	Graphical user Interface
HDF	Hierarchical Data Format
JDK	Java Development Kit
JRE	Java Runtime Environment
MGD	Multi-look Ground range Detected
NIC	U.S. National Ice Center
R&D	Research and Development
SAR	Synthetic Aperture Radar
SLC	Single Look Complex
SUMO	Search for Unidentified Maritime Objects
VDS	Vessel Detection System
UTC	Universal Time Coordinated

Annex. Q&A

Q: What is SUMO?

SUMO is a software package to perform ship detection on satellite SAR (Synthetic Aperture Radar) images. It has been developed at the Joint Research Centre of the European Commission.

Q: Which language is SUMO written in?

Java.

Q: On which operating systems does SUMO run?

Microsoft Windows and Linux.

Q: Under which license is SUMO being released?

SUMO is being released under the FreeBSD two-clause license.

Q: If I use SUMO in scientific research that results in a publication, how should I acknowledge it?

If you find SUMO useful, please cite the following paper:

H. Greidanus, M. Alvarez, C. Santamaria, F.-X. Thoorens, N. Kourti, P. Argentieri, 'The SUMO Ship Detector Algorithm for Satellite Radar Images', *Remote Sensing* **2017**, 9(3), 246; doi:10.3390/rs9030246

Q: In which modes can SUMO be run?

SUMO can be run in GUI and command line modes. In GUI mode, SUMO presents a visualisation of the image and detected targets and allows the user to manually inspect and correct the detections. The command line mode allows the user to easily batch-process a large amount of images.

Q: Which image types can be analysed by SUMO?

Please see section 4.1.1 (Image input data) for a list of satellite/sensors that are accepted. SUMO can only read the formats (unzipped) in which the products are typically delivered by the image providers.

Q: Which processing levels are accepted?

Generally, Ground Range Detected products are accepted. The Single Look Complex products of some sensors are also accepted. Geocoded products (also known as geocorrected or terrain-corrected) are not recommended: although they may be accepted by SUMO, some parts of the ship detection process may not perform as expected.

Q: Can SUMO detect ships in satellite optical images or airborne radar images?

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

However, limitations in terms of the image types the SUMO software package accepts still remain.

Q: Can SUMO detect oil platforms, offshore wind turbines, fishing gear, aquaculture equipment, icebergs or other non-ship objects?

Yes, similarly to ships, those structures will be detected if they pass the CFAR detection process. However, SUMO does not attempt to discriminate ships from non-ship objects.

Q: Does SUMO detect oil spills?

No, it does not.

Q: Does SUMO provide an estimate of the ships' speeds?

No, it does not.

Q: Does SUMO flag azimuth ambiguities? And range ambiguities?

SUMO detects ambiguities if they pass the CFAR detection process. Azimuth ambiguities are flagged with the value 'AA' in the <false_alarm_cause> tag of the output xml file. Greidanus et al., 2017 explains how this is done. SUMO does not flag range ambiguities.

GETTING IN TOUCH WITH THE EU

In person

All over the European Union there are hundreds of Europe Direct information centres. You can find the address of the centre nearest you at: <http://europa.eu/contact>

On the phone or by email

Europe Direct is a service that answers your questions about the European Union. You can contact this service:

- by freephone: 00 800 6 7 8 9 10 11 (certain operators may charge for these calls),
- at the following standard number: +32 22999696, or
- by electronic mail via: <http://europa.eu/contact>

FINDING INFORMATION ABOUT THE EU

Online

Information about the European Union in all the official languages of the EU is available on the Europa website at: <http://europa.eu>

EU publications

You can download or order free and priced EU publications from EU Bookshop at: <http://bookshop.europa.eu>. Multiple copies of free publications may be obtained by contacting Europe Direct or your local information centre (see <http://europa.eu/contact>).

JRC Mission

As the science and knowledge service of the European Commission, the Joint Research Centre's mission is to support EU policies with independent evidence throughout the whole policy cycle.



EU Science Hub
ec.europa.eu/jrc



@EU_ScienceHub



EU Science Hub - Joint Research Centre



Joint Research Centre



EU Science Hub



Publications Office

doi:10.2760/349278

ISBN 978-92-79-71881-6