

Monitoring of Oil Tank Filling with Spaceborne SAR using Coherent Scatterers

Carlos Villamil Lopez, and Uwe Stilla, *Senior Member, IEEE*

Abstract—In this paper a new method is introduced for the automatic estimation of all the relevant parameters of oil storage tanks using a single high resolution SAR image. For a given storage tank, this method will estimate its maximum capacity and determine whether it has a fixed or a floating roof. For tanks with a floating roof, the amount of oil stored will also be estimated. If a SAR time series is available all the images can be processed jointly, exploiting the available temporal information to provide more accurate and robust estimates than those obtained from each individual image. The dimensions of each storage tank are derived from its semicircular double reflections, which are detected using the coherent scatterers in the SAR image. The classification between tanks with fixed and floating roofs is performed with a simple machine learning classifier using just three features related to the detected semicircular double reflections. The performance of the proposed method using a single image and a time series is evaluated with three TerraSAR-X images of the port of Rotterdam, containing 167 oil storage tanks of different sizes and with both fixed and floating roofs.

Index Terms—SAR, synthetic aperture radar, time series, oil storage monitoring.

I. INTRODUCTION

CYLINDRICAL tanks are commonly used to store large quantities of petroleum products such as crude oil. Most of these oil storage tanks are located above ground [1], and are built following standards such as [2]. Two types of oil storage tanks can be distinguished: those with a fixed or a floating roof. Floating roofs rise and fall with the amount of fluid inside the tanks in order to decrease the vapor space above the liquid level, and are preferred for the storage of highly volatile fluids [1]. High resolution satellite images can be used to measure the dimensions (radius and height) of these cylindrical tanks and, if present, the position of the floating roof. This can be exploited to determine the maximum capacity of a given tank and, if the tank has a floating roof, also the level of fluid inside it. Therefore, satellite imagery can provide much more relevant information for oil tanks with a floating roof: by using a series of images acquired at different times, the oil production of a refinery with floating roof tanks can be effectively monitored. Because of the relevance that oil production has in the economy, this has recently become a popular commercial application of satellite data [3], [4].

C. Villamil Lopez is with the German Aerospace Center (DLR), Microwaves and Radar Institute, 82234 Wessling, Germany, and also with the Department of Photogrammetry and Remote Sensing, Technical University of Munich, 80333 Munich, Germany (e-mail: carlos.villamillopez@dlr.de).

U. Stilla is with the Department of Photogrammetry and Remote Sensing, Technical University of Munich, 80333 Munich, Germany (e-mail: stilla@tum.de).




Fig. 1. Appearance of an oil storage tank with a floating roof in a nadir optical image (a) and an oblique optical image (b), both obtained from Google Earth, and in a SAR image (c) acquired with TerraSAR-X.




Fig. 2. Appearance of an oil storage tank with a fixed roof in a nadir optical image (a) and an oblique optical image (b), both obtained from Google Earth, and in a SAR image (c) acquired with TerraSAR-X.

Two different types of satellite images are currently employed for the monitoring of oil storage: those acquired by optical and synthetic aperture radar (SAR) sensors. The appearance of an oil storage tank with a floating roof in both optical and SAR images can be seen in Fig. 1. For comparison, a tank with a fixed roof is shown in Fig. 2. When using optical images acquired close to nadir angle, such as the one shown in Fig. 1a, the position of a tank's floating roof can be estimated from the size of the shadow created by the tank's wall over this roof (since the sun's position at the time of the image acquisition is known). On the other hand, in off-nadir (i.e., oblique) optical images and SAR images, the position of the floating roof can be directly seen, as shown in Fig. 1b and Fig. 1c, respectively. Specifically, in SAR images the vertical displacement of the floating roof with respect to the bottom of the tank can then be directly measured by exploiting the well-known SAR layover effect, as it will be described later in detail in section II.

During the acquisition of the optical images, sunlight illumination and a cloud-free sky are required. Because the satellites acquiring these images are on low Earth orbits, each of them can only image a given location on Earth at a few

specific instants every several days. If the required sunlight illumination or atmospheric conditions are not given during these time slots, the monitoring of the oil storage tanks is not possible. This severely limits the applicability of such methods, especially for some locations where clouds and smog are often present. In the literature, most of the publications dealing with oil storage tanks in optical remote sensing images focus exclusively on the detection of such tanks [5]–[9], whereas just one publication has been found which describes an automatic method for the estimation of oil storage from the shadow in the floating roof [10].

Unlike optical sensors, SAR sensors can acquire images day-and-night, and those in X-band and lower frequency bands are practically not affected by clouds, fog or smog. Therefore, regular measurements of oil production can be ensured. The achievable accuracy depends on the sensor's spatial resolution, the chosen imaging geometry and the quality of the algorithm used to extract this information. However, this last point represents a challenge: due to radar specific imaging effects such as speckle noise, a robust and automatic extraction of this information from SAR images is far more challenging than in the case of optical images.

In the literature, several methods have been presented for the detection of oil storage tanks in SAR images [11]–[13], and there are four publications dealing with the estimation of the dimensions and the amount of oil stored in such tanks [14]–[17]. In [14] the authors show that it is possible to accurately estimate the height and diameter of cylindrical oil storage tanks with fixed roofs from SAR images, and that this can be achieved by exploiting the imaging geometry as well as its radiometry. They also mention that the same principle could be applied to estimate the fill level of oil storage tanks with a floating roof. In [15], a detailed analysis of the SAR signature of oil storage tanks with both fixed and floating roofs has been performed using SAR simulation, and the authors have additionally shown how all the relevant dimensions of the storage tanks can be estimated by measuring the distance between a few points on their SAR signature and taking into account geometric effects such as layover. The accuracy of these measurements has been verified using a TerraSAR-X image and ground truth data about the dimensions and fill level of the storage tanks, which was provided by the company operating the imaged oil tanks. An important limitation of the approaches presented in these two papers [14], [15] is that these measurements must be performed by a human operator, which has to manually select some specific image points for each oil storage tank. This greatly limits the applicability of these approaches for regularly monitoring the fill level of a large number of oil storage tanks.

More recently, two publications [16], [17] have introduced methods capable of automatically extracting some information about oil storage from SAR images. The method described in [16] can automatically estimate the vertical displacement over time of the floating roof of oil storage tanks using time series of moderate resolution SAR images. For this, a reference point is automatically selected for each storage tank by analyzing the minimum and average amplitude of the time series, and the location of the amplitude peak corresponding

to the double reflection at the tank's floating roof is detected for each individual image. The vertical displacement of the floating roof with respect to the reference point can then be obtained from the shift of this amplitude peak along the range axis. In [17], the authors present a method capable of estimating the dimensions and the amount of oil stored in tanks with a floating roof using a single high resolution SAR image. For this the image patch containing an oil storage tank is first denoised, and then some specific points of the storage tank are detected using two one dimensional amplitude profiles along the azimuth and range axes. The dimensions of the storage tank are then derived from these points. However, a description of the algorithm applied to automatically detect the points used to estimate the tank radius is not provided. The method was tested using KOMPSAT-5 SAR data, and the obtained results were compared with those estimated from optical images acquired by KOMPSAT-3.

In conclusion, all the methods employing optical images are affected by external factors (e.g., sun and clouds), and regular observations with optimal accuracy cannot be guaranteed. These limitations are imposed by inherent characteristics of optical sensors that cannot be easily changed. On the other hand, when using SAR images for monitoring the fill level of oil tanks, the main limitations are not imposed by the sensor, but rather by the accuracy and robustness of the available algorithms that can be employed for extracting this information. The complexity of automatic information extraction from SAR images can be seen by the fact that from the four methods that have been presented in the literature which deal with the estimation of oil storage, only the methods presented in [16], [17] are capable of extracting some of this information automatically without the need of manual interaction. The method introduced in [17] will be considered as the state-of-the-art, as it is the most recent publication and also capable of automatically extracting more information on oil storage.

In this paper a new method for the automatic and precise monitoring of oil storage using high resolution SAR images is presented, which has some significant advantages when compared to the current state-of-the-art:

- It can be applied for both storage tanks with floating and fixed roofs, and can automatically classify tanks according to their roof type. The method described in [17] only works with tanks with a floating roof.
- It exploits the complete geometry of the storage tanks by taking into account all the points in their semicircular double reflections, rather than just using a few select image points as in [17]. This increases the robustness and the accuracy that can be achieved.
- If multiple images are available, the proposed method can exploit them jointly to increase the accuracy of each measurement and the overall robustness of the method, rather than simply performing multiple individual measurements.
- It can track the vertical displacement of all the point scatterers in a floating roof with subpixel accuracy, which allows the precise estimation of changes in oil storage over time.

This paper is organized as follows: first, in section II the SAR signatures of oil storage tanks with floating and fixed roofs will be briefly analyzed, and we will show how their relevant parameters can be derived from the semicircular double reflections in them. Additionally, some constraints which simplify the detection of these semicircular double reflections will be established. Then we will describe how this can be applied to automatically extract all the relevant information for a given storage tank. This will be shown for two different cases: one in which a single SAR image is available, which will be described in section III, and another one in which a time series of SAR images is available, presented in section IV. Finally, results are shown in section V, and the conclusions are summarized in section VI.

II. SAR SIGNATURE OF OIL STORAGE TANKS

Oil storage tanks have a simple and well-known geometry and they exhibit a very characteristic SAR signature, which has been analyzed in detail via SAR simulation in [15] for both tanks with fixed and floating roofs. Multiple semicircular double reflections are present on the SAR signatures of both types of tanks. Knowledge of these semicircles is sufficient to estimate all the relevant parameters of an oil storage tank: its precise location and maximum capacity, and also its current fill level if the tank has a floating roof. Therefore, if these semicircles can be detected in the SAR image, they can be exploited to extract precise information about oil storage.

In this section we will analyze the geometry of both types of storage tanks, as well their temporal behavior and backscattering properties, and establish some constraints which can be enforced to simplify the detection of these semicircles. As the monitoring of oil storage tanks with a floating roof is the most interesting use case, their SAR signature will be analyzed first and in more detail. Then, the signature of an oil tank with a fixed roof will be briefly analyzed by comparing it to the signature of a tank with a floating roof.

A. Oil storage tanks with a floating roof

A SAR image showing an oil storage tank with a floating roof can be seen in Fig. 3a. This SAR image has been acquired by TerraSAR-X with an incidence angle of 48.1° , a resolution of 58 cm in slant-range and 23 cm in azimuth. For this visualization the original single-look complex (SLC) image has been resized to a square pixel spacing (in slant-range) and rotated so that the range axis corresponds to the image's y-axis, making the layover effect easier to interpret. This transformation will be applied to visualize all the results shown in this paper, but the actual processing will be always performed using the original SLC image rasters. As it can be seen in Fig. 3a, the SAR signature of an oil storage tank with a floating roof contains three bright semicircles, which appear as semi-ellipses in the image due to the chosen pixel spacing (i.e., corresponding to a slant-range projection). These are highlighted in Fig. 3b, and correspond to the three double reflections illustrated in Fig. 3c:

- One occurring between the outer wall of the cylindrical tank and the ground (shown in blue in Fig. 3b and 3c).

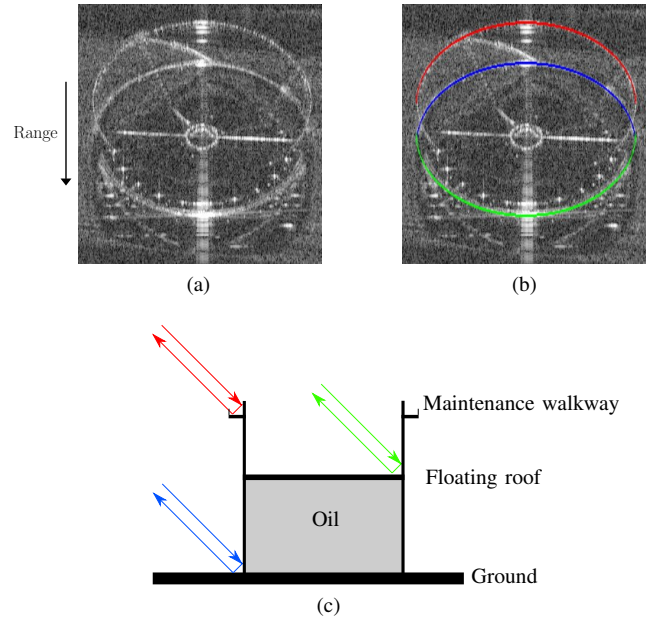


Fig. 3. The SAR signature of oil storage tanks with a floating roof contains three semicircular double reflections, as shown in: amplitude SAR image (a), amplitude SAR image with highlighted semicircular double reflections (b), and illustrative drawing of the geometry (c).

- Another between the outer wall of the tank and a maintenance walkway located near the top (shown in red in Fig. 3b and 3c).
- The last one between the inner wall of the tank and the upper surface of the floating roof (shown in green in Fig. 3b and 3c).

Even though these reflections typically appear as semi-ellipses in SAR images, we will refer to them as semicircles, as this is their true shape. The semi-axes of the ellipses in the SAR image can be easily calculated from the radius of the oil tank r_t , the mean incidence angle θ with which the corresponding SAR image was acquired, and the pixel spacing along the x axis (δ_x) and y axis (δ_y). Without loss of generality, we will assume throughout this paper that the x and y image axes correspond to the azimuth and range axes of the SAR image, respectively. The semi-axes of this ellipse r_{tx} and r_{ty} , expressed in image pixels, can be computed with the following equations:

$$r_{tx} = \frac{r_t}{\delta_x} \quad (1)$$

$$r_{ty} = \frac{r_t}{\delta_y} \sin \theta \quad (2)$$

When attempting to detect these semicircles in a SAR image, the following constraints, which are imposed by the geometry of the problem, can be enforced:

- The three semicircles have the same radius, as we can assume the thickness of the tank's wall to be negligible.
- Their centers lie along the same vertical axis, located at different heights. Therefore, the three semicircles will have the same center along the azimuth axis, but they will exhibit a translation in range due to the layover effect.

- Due to the side-looking imaging geometry of SAR sensors, only the side of the tank's outer wall which is closer to the SAR sensor can be imaged. Therefore, the two semicircles corresponding to the double reflections with the outer tank wall will appear towards near-range. These are highlighted in red and blue color in Fig. 3.
- Regarding the tank's inner wall, only the side which is farther away from the SAR sensor can be seen in the SAR image. Therefore, the double reflection corresponding to the inner wall of the tank (highlighted in green in Fig. 3) will appear towards far-range.

The aforementioned displacement of the semicircles centers along the range axis due to layover can be exploited to compute the height of the cylindrical tank h_t and the vertical position of the floating roof h_r by using the following equations:

$$h_t = \frac{l_t \delta_y}{\cos \theta} = \frac{(y_b - y_t) \delta_y}{\cos \theta} \quad (3)$$

$$h_r = \frac{l_r \delta_y}{\cos \theta} = \frac{(y_b - y_r) \delta_y}{\cos \theta} \quad (4)$$

where l_t and l_r represent the layover (in pixels) due to the height of the cylindrical tank and the vertical position of the floating roof, respectively. These can in turn be expressed as a function of y_b , y_t and y_r , which are the pixel coordinates along the y axis (range) of the centers of the semicircular double reflections at the tank bottom, top, and floating roof, respectively. Both equations assume that the range axis of the SAR image starts at near-range (such as in all the images shown in this paper), and will result in negative height values if the SAR image has an inverted range axis (i.e., starting at far-range), as in that case the layover effect will occur towards the opposite direction.

It is important to note that the height of the tank h_t might be slightly underestimated, as it is obtained from the double reflection occurring at a maintenance walkaway, which is slightly below the top of the tank. This error in the height of the cylindrical tank will only affect the estimation of its maximum capacity and not of the amount of oil stored inside it, which is the most interesting parameter. Additionally, this slight underestimation of a tank's height should be easy to compensate: if ground truth data is available for a few storage tanks, it should be possible to derive a offset and/or a scale factor to calibrate the estimated heights, as most oil storage tanks have standard sizes.

If a time series of SAR images acquired at different times is available, additional constraints associated with the temporal behavior can be imposed:

- The double reflections which correspond to the oil tank's outer-wall (red and blue semicircles in Fig. 3b) will remain constant over time, as the tank's outer structure does not change.
- On the other hand, when observed at different times the double reflection associated to the oil tank's inner-wall (green semicircle in Fig. 3b) can move along the range axis, as the floating roof will rise or sink with changes in the amount of oil stored.




Fig. 4. SAR images of an oil storage tank with a floating roof at two different dates: first image (a), second image (b), and multitemporal color composite image with the two amplitude images in the red and green channels, and coherence in blue.




Fig. 5. SAR amplitude image of an oil storage tank with a floating roof (a) and detected coherent scatterers (b).

This is illustrated in Fig. 4 for a pair of SAR images of the same oil storage tank acquired 11 days apart. The two individual amplitude images are shown in Fig. 4a and 4b. A multitemporal color composite image with both amplitude images in the red and green channels, and the interferometric coherence in the blue channel is shown in Fig. 4c. Here, it can be clearly seen how the double reflection corresponding to the floating roof moves between both image acquisitions, which appear in green and red colors in Fig. 4c due to the strong changes in SAR amplitude and the low coherence. On the other hand the two double reflections corresponding to the outer tank structure remain unchanged and appear therefore in blue/white colors, due to the high interferometric coherence and lack of change in amplitude. The outer tank structure is expected to remain unchanged and have high coherence even during long time periods. However, there is no guarantee that the floating roof will move, especially for short time periods, and could also appear unchanged. Therefore, both possibilities (the floating roof moving or remaining static) must be taken into account when exploiting the temporal information in a SAR time series for the monitoring of oil storage.

Finally, in addition to these geometric and temporal constraints, the backscattering properties can also be exploited for simplifying the detection of these semicircles. Many of the pixels in these semicircular double reflections appear in high resolution SAR images as point scatterers which dominate their corresponding resolution cells, which are sometimes denoted in the literature as coherent scatterers (CSs) [18]. By detecting these CSs in the SAR images, the detection of these semicircles can be greatly simplified. The reason for this is that the majority of the surroundings of the oil storage tank will be eliminated, leaving only a few relevant point scatterers

which can be used to fit the semicircles. An example of the CSs detected for an oil storage tank with a floating roof is shown in Fig. 5. The method used for the detection of CSs will be later described in detail in section III-A.

B. Oil storage tanks with a fixed roof

A SAR image of an oil storage tank with a fixed roof can be seen in Fig. 6a. This image and the one of a tank with a floating roof shown in Fig. 3a are part of the same TerraSAR-X image, and have therefore the same resolution and imaging geometry. By comparing both images it can be seen that the SAR signatures of both types of storage tanks are relatively similar, with the exception of the features corresponding to the floating roof. The same two semicircular double reflections at the bottom and top of the cylindrical tank are present, whereas the one corresponding to the floating roof is missing. This implies that the geometrical constraints established in the previous point regarding those two semicircles also apply to storage tanks with a fixed roof, and their height can also be computed using (3).

When observing a fixed roof tank at different times there will be no significant change, as no floating roof is present, and the two semicircular double reflections corresponding to the outer tank structure will remain unchanged. This can be seen in Fig. 6b, which is a multitemporal color composite image highlighting the changes between a pair of images acquired 11 days apart, such as the one shown previously in Fig. 4c for a tank with a floating roof. The only changes which can be seen are small amplitude changes in the clutter areas due to speckle noise. The pixels corresponding to the oil storage tank have a high interferometric coherence and no significant amplitude change, and appear therefore in blue and white colors.

The backscattering properties are also similar to those of a tank with a floating roof, and the pixels on the semicircular double reflections are also CSs. As in the previous case, the detection of these point scatterers allows to eliminate the tank's surroundings and can simplify the fitting of these semicircles. The detected CSs for this storage tank with a fixed roof are shown in Fig. 6c. Here, it can be seen that fewer CSs are detected at the double semicircular reflection at the top of the tank than in the case of a tank with a floating roof. However, there are still enough point scatterers to fit the corresponding semicircle and estimate the tank height using (3).

All these similarities can be exploited to develop a method capable of extracting all the relevant parameters of both types of tanks. Additionally, the main difference, which is the lack of the third semicircular double reflection corresponding to the floating roof, can be exploited to easily distinguish both types of tanks. Such a method, capable of classifying both types of tanks and extracting all their relevant parameters out of a high resolution SAR image, will be introduced in section III. Then, in section IV, this method will be extended to the case in which a SAR time series is available, exploiting the temporal information to achieve more accurate and robust estimates than those possible using each individual image.




Fig. 6. SAR amplitude image of an oil storage tank with a fixed roof (a), multitemporal color composite image showing temporal change across an image pair acquired 11 days apart (b), and detected coherent scatterers (c).

III. OIL STORAGE ESTIMATION FROM ONE SAR IMAGE

In the previous section it has been established that if the semicircular double reflections of an oil storage tank can be extracted from a SAR image, all its relevant parameters can be derived from them. In this section, a method which exploits the detection of coherent scatterers (CSs) and the geometric constraints described in the previous section to automatically detect these semicircular double reflections will be presented. After these double reflections have been detected for a given storage tank, a simple machine learning classifier can be used to determine whether this tank has a floating or a fixed roof. Given one high resolution SAR image of a oil refinery, the proposed method can be applied to automatically estimate the maximum capacity of each storage tank, as well as the amount of oil stored in the tanks with a floating roof.

A. Pre-processing: coherent scatterer detection

Strong point scatterers, also known as coherent scatterers (CSs), can be detected in a high-resolution single-look complex (SLC) SAR image by applying spectral diversity techniques [18]. Multiple methods have been presented in the literature for the detection of CSs [18]–[21]. All of these methods are based on the same principle: they exploit the high bandwidth available in high resolution images by computing multiple sublook images with lower spatial resolution, and then detect CSs by identifying those targets which remain stable across these sublooks. The main difference between these methods is the criterion used to determine which targets remain stable across the multiple sublooks. A comparison of the performance of the different methods can be found in [18]. In this paper, the phase variance approach (PVA) [20] will be used for CS detection, due to its simplicity and relatively good performance. This method will be briefly described below.

The basic idea behind the CS detection using PVA is that if a given image pixel contains a CS its phase will vary linearly with sublook frequency, whereas if not it will vary randomly. Therefore, multiple sublooks with a reduced bandwidth and different central frequencies should be computed, so that the phase trend can be analyzed for each pixel. A description of the process of sublook calculation from a single-look-complex SAR image can be found in [18]. In this paper, the sublooks will be computed along range (i.e., corresponding to different frequency sub-bands), as is typically done for CS detection [18]–[21]. The same detection method could also be applied

to azimuth sublooks [18], but the performance is expected to be worse in that case, due to the fact that many point scatterers have a nonconstant azimuth angular scattering pattern [19].

For a given pixel, we can denote its phase at each sublook as ϕ_i , with $i = 1, \dots, n$, where n is the number of sublooks. To avoid issues when analyzing the phase trend due to phase jumps of 2π around $\pm\pi$, phase unwrapping should be performed by adding or subtracting multiples of 2π as required so that $|\phi_{i+1} - \phi_i| \leq \pi$. To determine whether the pixel's phase varies linearly with frequency, the mean and variance of the phase derivative with respect to frequency can be computed:

$$\mu_\phi = \frac{1}{n-1} \sum_{i=1}^{n-1} \frac{\phi_{i+1} - \phi_i}{f_s} \quad (5)$$

$$\sigma_\phi^2 = \frac{1}{n-1} \sum_{i=1}^{n-1} \left(\frac{\phi_{i+1} - \phi_i}{f_s} - \mu_\phi \right)^2 \quad (6)$$

where f_s is the difference between the central frequencies of two consecutive sublooks, given in MHz. If this pixel has a linear phase, the phase derivative should be nearly constant and therefore have a low variance σ_ϕ^2 . Therefore, CSs can be detected by computing σ_ϕ^2 for each pixel and applying a threshold, which can be denoted as T : pixels with $\sigma_\phi^2 < T$ will be considered to contain a CS.

For the pixels with a linear phase, the slope of this phase ramp will be directly related to the distance along the range axis between the center of that pixel and the actual location of the corresponding CS [18]. If μ_ϕ is used as a estimation of the slope of this linear phase ramp, this distance, which can be denoted as d_r , can be expressed as:

$$d_r = \frac{c}{4\pi} \mu_\phi \quad (7)$$

where c is the speed of light. This relation can be exploited to determine the location of CSs along the range axis with a higher accuracy than that given by the range resolution of the SAR sensor. Additionally, it can also be used to mitigate the resolution loss introduced by sublook computation, which will cause strong point scatterers to spread to neighboring pixels. In such cases, those neighboring pixels might also have a linear phase trend even if they do not contain a CS, and would be falsely identified as CSs when just applying a threshold to σ_ϕ^2 . However, because their phase ramps will have a steeper slope, these can be filtered out by also thresholding μ_ϕ . The corresponding limits for μ_ϕ can be analytically derived by enforcing that a CS must be inside the resolution cell of a pixel (i.e., less than half a pixel distance between the CS and the pixel center): $|d_r| \leq 0.5 \delta_y \Rightarrow |\mu_\phi| \leq \frac{2\pi\delta_y}{c}$.

For the detection of CS in this paper the following parameters were used: a threshold $T = 0.005$, and $n = 40$ sublooks covering the whole available bandwidth with a 75% spectral overlap between consecutive sublooks. For the TerraSAR-X imagery used in this paper with 300 MHz total bandwidth, this resulted on a sublook bandwidth of 27.90 MHz and a separation of $f_s = 6.97$ MHz between sublooks. An example of the results obtained using the described method can be seen in Fig. 5.

B. Approximate location of the oil storage tanks

Before attempting to detect the semicircular double reflections of a given storage tank, its approximate location in the SAR image should be known. As the goal of the method presented in this paper is not the detection of oil storage tanks, but rather the automatic and precise estimation of their dimensions and of the amount of oil stored in the tanks with a floating roof, we will assume that the approximate locations of the oil tanks are known. Here, a brief overview of the different methods and data sources which can be used for obtaining the locations of the storage tanks will be provided.

One possibility would be to obtain these locations from GIS data, such as OpenStreetMap (OSM), which is open and freely available and contains the locations of many of such oil tanks. For the oil refineries which are not present in OpenStreetMap, the locations of their storage tanks could also be detected in the SAR images using any of the approaches that have already been presented in the literature for the detection of cylindrical storage tanks [11]–[13]. Alternatively, the oil tanks could be detected also in an optical image if one is available, using one of the methods described in [5]–[8]. If enough optical imagery is available, a convolutional neural network (CNN) could also be trained to detect the oil tanks, as described in [9], [22]. For this, a large training dataset can be semi-automatically generated by using OpenStreetMap data to obtain the location of many oil storage tanks, generating in this way the labels for the optical imagery. In addition to the approximate location of the storage tanks, most of the listed methods also provide a bounding box around each tank, which can be used to estimate an approximate value for its radius. While the method proposed in this paper does not require an initial estimate for the radius of a given tank, it can be used if available to reduce the computing time and slightly increase the robustness of the method, as it will be shown later.

For the examples shown in this paper, the approximate locations of the oil tanks were obtained automatically from OpenStreetMap by performing a query using its Overpass API. For a given storage tank, the latitude and longitude coordinates of its center obtained from OpenStreetMap are projected into the SAR image by performing geocoding with a digital elevation model (DEM), obtaining the approximate pixel position for the center of the tank bottom $\hat{\mathbf{p}}_b = (\hat{x}_b, \hat{y}_b)$. Additionally, an approximate value for the radius \hat{r}_t of each tank was also estimated from OpenStreetMap data. For this, the polygon associated to each tank first projected into a UTM map projection, in order to avoid errors in the estimated radius due to the map projection used by OpenStreetMap.

Independently on the method used to obtain these locations, this step only needs to be performed once, as the locations of such oil tanks change very rarely, and the precise location and other parameters will then be extracted from the SAR image using the method described below.

C. Estimation of the precise location and size of oil tanks

Once the CSs in the SAR image have been detected and the approximate locations of the storage tanks in the SAR image are known, the semicircular double reflections of each tank

should be detected to obtain information about oil storage. For a given storage tank, the goal is to determine the following parameters: its radius r_t , height h_t , and vertical position of the floating roof h_r (if there is one). Additionally, the exact pixel position for the center of the tank bottom $\mathbf{p}_b = (x_b, y_b)$ will also be determined, from which the corresponding geographic coordinates can easily be computed by applying geocoding. Out of these parameters, the maximum capacity of the storage tank V_{max} (given in cubic meters) can be easily computed using (8). If the tank has a floating roof, the volume of oil V_{oil} stored inside it at the time the SAR image was acquired can also be computed in the same way:

$$V_{max} = \pi r_t^2 h_t \quad (8)$$

$$V_{oil} = \pi r_t^2 h_r \quad (9)$$

Both equations assume that the volume available for oil storage begins at the height of the lower semicircular double reflection, which is expected to occur at ground height. This assumption should be valid for aboveground storage tanks (the most common type of tanks) [1], as the thickness of the tank's bottom plate is expected to be negligible [2]. However, a bias could possibly be introduced in some cases by a foundation of unknown height below the tank: in [15], the authors observed a bias in the tank fillings estimated from a SAR image, which they attributed to a foundation height of approximately 1 m.

For the tank radius r_t and height h_t minimum and maximum possible values can be specified, taking into account the possible sizes for these storage tanks. For the examples shown in this paper, the following limits for the height h_t and radius r_t have been chosen, which covered the sizes of all the storage tanks present in the available SAR imagery: $h_t^{min} = 12.5$ m, $h_t^{max} = 25$ m, $r_t^{min} = 10$ m and $r_t^{max} = 50$ m. These limits for the tank height h_t can be translated into the corresponding limits for the layover l_t (in pixels) using (3).

If the method previously used for obtaining the approximate location of the storage tanks also provides an approximate value for their radius \hat{r}_t , it can be used for setting tighter limits for r_t , specified individually for each tank. In section V, the results and runtimes will be compared for the cases in which the generic limits for r_t are used, with those obtained when an approximate radius \hat{r}_t is available for each tank.

Besides these size limits, the only actual input required is the approximate pixel position for the center of the tank bottom $\hat{\mathbf{p}}_b$, which should have been obtained using any of the methods discussed in the previous subsection. Additionally, an uncertainty value u giving the maximum expected error (in meters) of this approximate position should be specified. For the examples in this paper, it was assumed that the locations obtained from OpenStreetMap data had an uncertainty $u = 15$ m. This uncertainty u , can be scaled into pixels for both the x axis (u_x) and y axis (u_y) using the respective pixel spacings and incidence angle, in the same way the radius r_t was scaled in equations (1) and (2).

Using the approximate location of the storage tank $\hat{\mathbf{p}}_b$ and its associated uncertainty u , as well as the established limits for the tank radius r_t and layover l_t , a small image patch




Fig. 7. Detection of semicircular double reflections at the top and bottom of a storage tank: CSs inside the image patch surrounding a storage tank (a), possible center pixels illustrated for three different CSs and two different radii (b), and two semicircles detected using the Hough accumulator h_2 (c).

containing the oil storage tank can be obtained by setting the following boundaries for x and y :

$$\hat{x}_b - u_x - r_{tx}^{max} \leq x \leq \hat{x}_b + u_x + r_{tx}^{max} \quad (10)$$

$$\hat{y}_b - u_y - r_{ty}^{max} - l_t^{max} \leq y \leq \hat{y}_b + u_y + r_{ty}^{max} \quad (11)$$

The CSs detected inside this image patch can then be used to fit the semicircles corresponding to the double reflections of the storage tank. An example of such an image patch can be seen in Fig. 7a, which shows the detected CSs for a storage tank with a floating roof. It is important to note that even though this and all the images shown in this paper have a square pixel spacing in slant range and the range direction along the y -axis, this transformation is only applied for visualization purposes and the actual processing is done using the original SLC SAR image raster.

Initially, the two semicircles at the bottom and at the top of the outer tank structure (highlighted in blue and red colors in Fig. 3) should be detected, as these are present for both the storage tanks with a fixed and a floating roof. Both semicircles appear towards near-range (i.e., towards the top of the y axis in the shown examples), and are defined by the parameters x_b , y_b , r_t and l_t : both have a radius r_t (given in meters), the lower semicircle has its center at pixel (x_b, y_b) , and the upper one has its center at pixel $(x_b, y_b - l_t)$. Their detection, and therefore the estimation of these four parameters, can be formulated as an optimization problem: the goal is to find the values for these parameters which maximize the number of CSs that fit the corresponding semicircles. While a brute-force search could be applied to test all the possible parameter values inside their respective intervals, a more efficient implementation is possible by using a Hough transform. Below, a brief explanation will be provided on how a Hough transform can be used for the detection of these two semicircles on the SAR image.

The Hough transform will compute a three dimensional accumulator array $h(x, y, r)$, whose values will give the number of CSs fitting the semicircle of radius r centered at (x, y) . All the elements of this accumulator array must be initially set to zero, and will be updated iteratively. For a given radius r_0 and a CS located at pixel \mathbf{p}_i , $h(x, y, r_0)$ should be incremented by 1 for all the pixels (x, y) which could be the center of a semicircle of radius r_0 passing through \mathbf{p}_i . Here, the fact that semicircles appear as semi-ellipses in the SAR image must be accounted for. This process should be repeated for all the

detected CSs, and for multiple values of the radius r covering the interval between r_t^{min} and r_t^{max} . Figure 7b illustrates how the computation of the Hough accumulator works, showing the possible centers for three different CSs (each highlighted in a different color) and two different radii (one shown as continuous and the other as dashed lines). For sampling the interval for the radius r , it is recommended to select a step δ_r which ensures that the corresponding change in image pixels is equal to or smaller than one pixel for both the x and y axes:

$$\delta_r = \min(\delta_x, \delta_y / \sin \theta) \quad (12)$$

The accumulator array $h(x, y, r)$ will tend to have lower values for smaller semicircles, even if they represent a good fit. This can be compensated by computing a new accumulator array $h'(x, y, r)$ which accounts for this:

$$h'(x, y, r) = h(x, y, r) / \sqrt{r} \quad (13)$$

Even though the number of CSs along each semicircle grows linearly with the radius r , the normalization was performed by dividing by a factor of \sqrt{r} instead, as this has shown in our experiments to provide the best results for both small and large semicircles. From $h'(x, y, r)$, a new accumulator array to detect two semicircles with a displacement of l pixels along the y axis can be easily defined:

$$h_2(x, y, r, l) = h'(x, y, r) + h'(x, y - l, r) \quad (14)$$

Finally, the values of x_b , y_b , r_t and l_t can be obtained from $h_2(x, y, r, l)$:

$$(x_b, y_b, r_t, l_t) = \arg \max_{\substack{x, y, r, l \\ x \in [\hat{x}_b - u_x, \hat{x}_b + u_x] \\ y \in [\hat{y}_b - u_y, \hat{y}_b + u_y] \\ r \in [r_t^{min}, r_t^{max}] \\ l \in [l_t^{min}, l_t^{max}]}} h_2(x, y, r, l) \quad (15)$$

For the example of Fig. 7, the two semicircles detected using (15) and the Hough accumulator h_2 can be seen in Fig. 7c.

At this point, the center pixel of the tank bottom and its radius are known, and the tank height h_t can then be easily computed from the obtained layover l_t using (3). This applies for both tanks with a floating or a fixed roof. However, for tanks with a floating roof, its vertical position h_r is still unknown.

D. Estimation of the floating roof position

The vertical position of the floating roof h_r can be easily computed from the corresponding layover l_r , which can be obtained by detecting the semicircle due to the double reflection of the floating roof, shown in green in Fig. 3. This semicircle also has a radius r_t , its center is at $(x_b, y_b - l_r)$ and, in contrast to the two previous semicircles, it appears towards far-range (i.e., towards the bottom of the y axis in the shown examples). As only one parameter (l_r) needs to be determined, there is no need to apply a Hough transform to detect this semicircle. Instead, cross-correlation with a semicircular binary mask can be performed, displacing it along the range axis (y) over all the possible center locations. For this, the minimum and maximum possible values of l_r need to be established, which




Fig. 8. Example for the estimation of the vertical position of the floating roof position for the storage tank previously shown in Fig. 7.

can be easily obtained from the corresponding limits for the floating roof height h_r . If a tank is completely empty, the floating roof will be at the bottom, and therefore $h_r^{min} = 0$. On the other hand, if it is full, the floating roof will be at the top, and $h_r^{max} \simeq h_t$. As previously mentioned, the tank height h_t is obtained from a double reflection at a maintenance walkway near the top of the tank, and is therefore slightly underestimated. Additionally, if the radius r_t obtained from the Hough transform is slightly larger than the actual radius of the storage tank, a small error will be induced in the estimated center y_b , which in turn will cause to slightly overestimate l_r . To account for this, it is recommended to set an upper limit for h_r which is larger than h_t . For the experiments shown in this paper, this limit was set to $h_r^{max} = h_t + 5$.

The cross-correlation between these two binary images will give the number of CSs which match each semicircle inside the defined height interval. The best fitting semicircle, and therefore the layover l_r (in pixels) due to the vertical position of the floating roof, can be obtained by finding the maximum of this cross-correlation. If we denote as $f(l)$ the function with the cross-correlation values for each layover displacement l , this can be formulated as it follows:

$$l_r = \arg \max_{l \in [0, l_r^{max}]} f(l) \quad (16)$$

where l_r^{max} is the maximum possible layover value, computed from h_r^{max} using (4). In Fig. 8, the results of this cross-correlation can be seen for the storage tank previously shown in Fig. 7. Here, the maximum can be clearly identified, and the estimated layover of the floating roof corresponds to the one that can be manually measured in the image.

If a storage tank is very tall and has a comparatively small radius, the floating roof might not be visible if it is located near the tank's bottom and the SAR image was acquired with a flat incidence angle, as it will be hidden by the radar shadow caused by the cylindrical outer wall. In order to be able to properly detect the semicircular reflection of the floating roof when this is located at the bottom, we can establish the following constraint: $\tan \theta \leq 1.5r/h_t$. As long as this condition is satisfied, an arc of around 80° of this semicircle should be visible. In real scenarios this is not expected to significantly

limit the applicability of the method, as this condition should be easily fulfilled for most tanks with conventional sizes: any incidence angle smaller than 50° would work for all the storage tanks in the dataset used in this paper (which is described later in section V). From this analysis it follows that using SAR images acquired with a steep incidence angle is advantageous for this method, as this allows to see better into the inside of the tank, and additionally it also allows to derive the heights from the layover with a higher accuracy.

A priori it is often unknown whether a tank has a fixed or a floating roof. In such cases the estimation of h_r can be performed as described, even if this only makes sense for tanks with a floating roof. If a given tank has a fixed roof, the cross-correlation value obtained when estimating h_r will be very low, as the double reflection of the floating roof is not present. This information can then be used to determine whether a tank has a fixed or a floating roof, as it will be described below.

E. Classification of storage tank type

As described in section II-B, the main difference between the SAR signature of a storage tank with a floating roof and one with a fixed roof is that in the latter there is no semicircular double reflection towards far-range (as it is caused by the floating roof). This will imply that when attempting to estimate the vertical position of the floating roof by detecting its double reflection (as done in the previous step), the number of matched CSs will vary significantly depending on whether the storage tank has a fixed or a floating roof: for a fixed roof this number will be very low, and for a floating roof it will be much higher. This number is given by $f(l_r)$, as defined in section III-D.

A further difference between both types of storage tanks is that for a tank with a fixed roof, less CSs are detected at the double reflection at its top than at the reflection at the bottom, as shown in Fig. 6c. On the other hand, for a tank with a floating roof the number of CSs detected on these two double reflections is similar, as it can be seen in Fig. 5b. The number of CSs at the bottom and top double reflections are given by $h(x_b, y_b, r_t)$ and $h(x_b, y_b - l_t, r_t)$ respectively, as defined in section III-C.

Taking all this into account, we propose to use the number of detected CSs at each of these three double reflections as the features to be fed into a machine learning classifier such as a Support Vector Machine (SVM) or a random forest. For each storage tank, the corresponding three dimensional feature vector \mathbf{v} can be defined as:

$$\mathbf{v} = (f(l_r), h(x_b, y_b, r_t), h(x_b, y_b - l_t, r_t)) \quad (17)$$

As we are only using three simple but informative features, it will be possible to train an accurate classifier with little training data. Additionally, as oil storage tanks have standard shapes and sizes, a classifier trained with storage tanks located at a given refinery should also be applicable to storage tanks at different locations. The classification results obtained by applying this approach to a dataset with over 150 storage tanks will be shown in section V-C. Here, different classifiers will also be evaluated, as well as the effect that increasing the amount of training samples has on the classification results.

IV. OIL STORAGE ESTIMATION FROM SAR TIME SERIES

In the previous section, a method to estimate all the relevant parameters of oil storage tanks out of a single high resolution SAR image has been introduced. An interesting application of such a method is the monitoring of changes in oil storage, by observing many storage tanks with a floating roof using series of SAR images acquired at different times. In such a scenario in which a SAR time series is available, the presented method can be improved by enforcing the temporal constraints established in section II. In this section we will show how the additional observations can be exploited to make the method more robust, enable more accurate estimates of several relevant parameters, and improve the classification between tanks with a fixed roof and those with a floating roof. To achieve this, a change detection method will be first introduced, which will then be applied to separate the coherent scatterers which remain static (e.g., those corresponding to the outer structure of a storage tank) from those that move over time (e.g., those corresponding to a floating roof). Then, the previously introduced method will be extended to take this temporal information into account.

A. Pre-processing: change detection with coherent scatterers

If two or more repeat pass acquisitions are available, coherent change detection can be applied to determine which CSs remain static and which move during the corresponding time period. Even though computing the interferometric coherence involves spatial averaging using a small window (which may contain clutter in addition to the CS), strong point scatterers tend to have high coherence values in high resolution SAR images and they do not exhibit significant temporal decorrelation [23], [24]. Therefore, as long as a CS remains static the corresponding pixel should have a high coherence, whereas even a small displacement (i.e., of less than one pixel) will cause a significant drop in the coherence. A change detection algorithm based on this will be described in detail below.

Given a time series of n single-look complex (SLC) SAR images, these should be first co-registered with subpixel accuracy by using a method such as the one described in [25]. After this, the CSs should be detected for each image in the stack as described in section III. This will result in a stack of n binary images \mathbf{I}_i , with $i = 1, \dots, n$, which will have a value of 1 for the pixels containing a CS and 0 for the rest. To detect which of these CSs have moved during the period covered by the time series, a threshold can be applied to the interferometric coherence. If the coherence drops below this threshold for a pixel containing a CS it will imply that this CS has moved, whereas if it is always above this threshold it will mean that it has remained static.

To implement this, the interferometric coherence must be computed for the $n - 1$ consecutive image pairs, obtaining a stack with the respective coherence images \mathbf{C}_j , with $j = 1, \dots, n - 1$. When computing the coherence a boxcar filter with a window of 7×7 pixels was used for the examples in this paper. Then the minimum coherence value can be computed for each pixel, obtaining another image \mathbf{C}_{min} . In the same way, the maximum of the binary images \mathbf{I}_i should

also be computed pixelwise, obtaining another binary image I_{max} showing all the pixels for which at least one CS has been detected in the time series. A binary image S with the CSs which remain static during the time series can then be computed by applying the following equation:

$$S(x, y) = \begin{cases} 1 & \text{if } C_{min}(x, y) > c_t \text{ and } I_{max}(x, y) = 1 \\ 0 & \text{otherwise} \end{cases} \quad (18)$$

where c_t is the chosen threshold for the interferometric coherence, which has been set to $c_t = 0.35$ for the examples shown in this paper. A further implication of (18) is that if a CS always has a coherence above c_t , it will be considered that it has remained static and it is therefore present in all the images, even if this CS was detected in just a single image of the series. The reasoning behind this is that a high coherence value should only be possible if there is no change, whereas false negatives during the CS detection are much more likely to occur.

Finally, by comparing the static CSs in S with those detected in each individual image I_i , a set of new images M_i showing the CSs that moved can be obtained:

$$M_i(x, y) = \begin{cases} 1 & \text{if } I_i(x, y) = 1 \text{ and } S(x, y) = 0 \\ 0 & \text{otherwise} \end{cases} \quad (19)$$

The results of this change detection algorithm can be seen in Fig. 9 for an image pair showing a storage tank with a floating roof at two different dates. Figure 9a shows this image pair as a multitemporal color composite image with both amplitude images in the red and green channels, and the interferometric coherence in the blue channel. The CSs detected in each individual image are shown in Fig. 9b and 9c. After the change detection, these CSs are classified into those which remained static (shown in Fig. 9d) and those that moved (shown for each image in Fig. 9e and 9f). As expected, it can be clearly seen that most of the static CSs correspond to the two double reflections at the outer tank structure, whereas most of the moving CSs correspond to the floating roof, which moved vertically due to a change in the amount of oil stored. It is important to note that for time series covering short time periods, the floating roof of some storage tanks might not move, as the amount of oil stored in them might remain constant during this time. In these cases, it will not be possible to separate the floating roof from the outer tank structure, as both will remain static. Because of this, the use of time series covering a long time period is advantageous.

B. Estimation of the precise location and size of oil tanks

Once the described change detection method has been applied to identify which CSs remained static and which moved, this temporal information can be exploited when estimating the relevant parameters of oil storage tanks. As in the previous scenario in which a single SAR image was available, initially the size and location of each storage tank should be determined. For this, the approximate locations of the storage tanks in the SAR images should be first obtained as described in section III-B. Then, a Hough transform can be

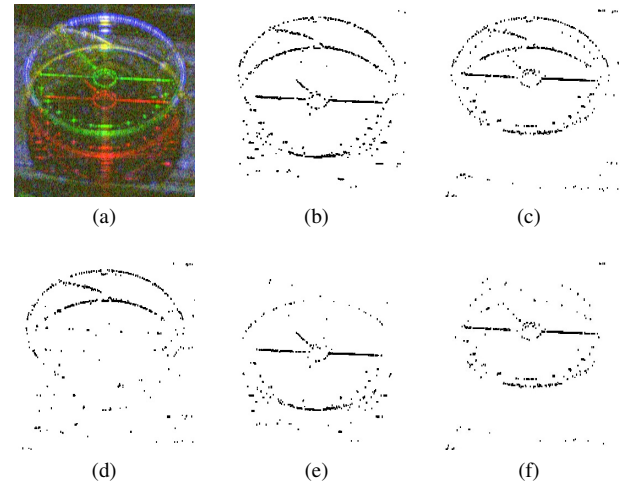


Fig. 9. Separation of static and moving parts of an oil storage tank using change detection with an image pair. Input image pair shown as a color composite image (a) and detected CSs for each image (b, c). Change detection results: static CSs (d) and moving CSs in each image (e, f).

applied as described in section III-C to estimate the parameters r_t (tank radius), h_t (its height), and p_b (the pixel position for the center of the tank bottom). In this case, only the CSs which remain static (i.e., those in the binary image S) should be taken into account when performing this Hough transform. For tanks with a floating roof which moved, this allows to separate the outer tank structure from the floating roof, which in turn makes the detection of the two semicircles faster and more robust, as most of the CSs which do not correspond to these semicircles can be filtered out.

For a given storage tank, once its precise location and size have been determined, what remains is to determine whether this tank has a fixed or a floating roof and, for tanks with a floating roof, the roof position at the acquisition time of each image must be estimated. Again, initially it will be assumed that all the storage tanks have a floating roof, and we will attempt to estimate its position. The results of this step will be used later to determine whether a storage tank has a floating or a fixed roof. Rather than applying the same approach as in the single image case to estimate the position of the floating roof, this can now be decomposed in two separate steps: the estimation of the vertical displacements of the floating roof between each successive image pair, and the estimation of its initial height.

C. Estimation of the vertical displacements of the floating roof

When the amount of oil stored in a tank with a floating roof changes, the floating roof will move up or down accordingly. Using a pair of SAR images acquired at different times, the corresponding vertical displacement of a given floating roof can be directly estimated, without the need to estimate its vertical position in each image.

The vertical displacement of a floating roof between the acquisition times of any two images i and j in the time series will be denoted as d_{ij} , and can be defined as:

$$d_{ij} = l_{r_j} - l_{r_i} \quad (20)$$




Fig. 10. Result from the estimation of the vertical displacement of the floating roof using cross-correlation with the image pair from Fig. 9.

where l_{rj} and l_{ri} denote the layover due to the vertical position of the floating roof at the images j and i in the series. Here the value of d_{ij} is given in pixels, but it can be easily converted into meters in the same way the layover l_r is scaled into the height h_r , using the relation defined in (4).

If a floating roof moves at some point during the period covered by the time series, its vertical displacement d_{ij} can be estimated using the binary images M_i and M_j , which contain the CSs which moved during the time series and are present at images i and j , respectively. Similar to the previously estimated parameters, the estimation of d_{ij} can also be seen as an optimization problem: the vertical displacement for which more CSs match between the two image patches of M_i and M_j will most likely correspond to the correct solution. This process will match all the features in the floating roof (not only the semicircular double reflection), and can be easily implemented by performing a cross-correlation of these two binary images for different displacements along the range axis (y). For this, a small image patch containing just the corresponding storage tank without any surroundings should be used. The corresponding image region can be obtained by modifying (10) and (11) using the tank location, radius and height values estimated in the previous step, as well as an uncertainty value $u = 0$. Combining (20) with the limits for l_r previously established in III, it follows that the value of d_{ij} must be within the following range: $-l_r^{max} \leq d_{ij} \leq l_r^{max}$. The maximum cross-correlation value obtained when estimating d_{ij} can be denoted as w_{ij} , and its value will be used at a later step.

The results obtained when applying this cross-correlation to the image pair from Fig. 9 can be seen in Fig. 10. The maximum of this cross-correlation can be clearly identified, and as expected the estimated vertical displacement corresponds to the one which can be manually measured.

If desired, the roof displacements can also be estimated with subpixel accuracy. As described in section III-A, the location of each detected CS can be determined along the range axis with subpixel accuracy. This can be exploited to compute a subpixel shift, which can be added to the shift obtained from the cross-correlation (which has an accuracy of one pixel). During the cross-correlation step, w_{ij} CSs

have been matched between the two images i and j for a displacement of d_{ij} pixels along range. The distance between each of these CSs and the corresponding pixel center can be determined separately for each image by applying (7). For each one of the matched CSs, the difference between the values obtained from the two images will give an estimate of its subpixel displacement. An accurate estimate of the subpixel displacement for the floating roof can be then obtained by averaging all the estimated subpixel displacements for the w_{ij} matched CSs.

For a time series with n images, the vertical displacement of a given floating roof between each of the $n - 1$ consecutive image pairs will provide all the required information on the changes in the amount of oil stored inside the tank. However, the described method can be used to estimate the roof's vertical displacement between each of the $\binom{n}{2} = (n^2 - n)/2$ unique image pairs in the series. The additional estimated values will provide some redundant information, as the displacements from multiple image pairs are related (e.g., the displacement between the first and third images must be the sum of the displacements between first and second images, and second and third images). This relation can be formulated in a general way as:

$$d_{ij} = d_{ik} + d_{kj} \quad (21)$$

For convenience, we will assume that the displacements are computed for the image pairs with $j > i$ (the remainder image pairs with $j < i$ do not provide any additional information). This redundant information can be exploited to obtain a more robust and accurate estimate of the changes in oil storage, as the displacement values estimated from each individual image pair using cross-correlation may contain small errors, and in some rare cases even a few of these estimates may be significantly off.

These improved estimates of the vertical displacement of the floating roof between each of the $n - 1$ consecutive image pairs will be denoted as d'_k , with $k = 1, \dots, n - 1$. Their values can be obtained by solving an overdetermined linear equation system with $(n^2 - n)/2$ equations (one for each estimated d_{ij}) and $n - 1$ unknowns (the values of d'_k to be estimated). To formulate this linear system, each of the estimated d_{ij} must be expressed as a linear combination of the unknowns, which can be done using the following equation:

$$d_{ij} = \sum_{k=i}^{j-1} d'_k \quad (22)$$

Such a overdetermined linear system can be solved by applying a method such as least squares. However, we propose to use the Huber regressor, which is a robust regression method and less sensitive to outliers than least squares. When performing this regression, the maximum cross-correlation values w_{ij} obtained when estimating each d_{ij} term can be used as weights, as typically those estimates with lower cross-correlation values will be less accurate.

Using this approach, the vertical displacements of a floating roof (and therefore the changes in oil storage) can be estimated with very high accuracy. The larger the number of SAR images in the series, the more robust and accurate that these estimates

will become. It is important to note that the estimated values will only be valid if the floating roof moved during the period covered by the time series. However, this is not an issue, as the case in which the floating roof does not move will also be accounted for next, when estimating its initial position.

D. Estimation of the initial position of the floating roof

Finally, in order to have an absolute measurement of the amount of oil stored in a given tank with a floating roof at the acquisition time of each image, the vertical position of the roof in the first image must be determined. If the floating roof moved at some point during the period covered by the time series, its vertical displacements have already been estimated in the previous step, and therefore only its initial position is missing. If the roof did not move, the displacements estimated in the previous step will be invalid, but the initial position of the floating roof will be all of the information that is required anyway. Here, a simple decision rule will be established to determine whether a floating roof has moved or not, and a method to estimate its initial position will be described for both of these scenarios.

If many CSs in the image patch containing a storage tank moved, this will imply that its floating roof has very likely moved. The weights computed in the previous step when estimating the vertical displacements of the roof provide a good indication for this: each weight w_{ij} quantifies the maximum number of moving CSs matching between the corresponding image pair. The average of all these weights, denoted as w_{avg} , can be used as a simple metric to quantify the number of CSs which could potentially correspond to a moving roof.

If the floating roof did not move, there will be very few CSs which moved, and the value of w_{avg} will be low. Additionally, the floating roof (and therefore its semicircular double reflection) will be present in the binary image \mathbf{S} which shows the CSs which remained static. Therefore, if many static CSs match a semicircle towards far-range, it will indicate that the roof did not move. To obtain the number of static CSs which match this semicircle, the method previously described in section III-D can be used. Rather than applying this method to the CSs detected in a single image, in this case it should be applied to the static CSs in \mathbf{S} . Now, the function with the cross-correlation values for each layover l can be denoted as $f_s(l)$, and the layover l_{r_s} for which the highest number of static CSs match such a semicircle can be obtained using equation (16).

The values of w_{avg} and $f_s(l_{r_s})$ can now be compared: if $f_s(l_{r_s}) > w_{avg}$ the floating roof is assumed to have remained static, and its layover will be given by l_{r_s} , which can be converted into the corresponding height using equation (4). On the other hand, if $w_{avg} \geq f_s(l_{r_s})$ the roof is assumed to have moved. The vertical displacements of the floating roof have been obtained in the previous step, but its initial position still needs to be determined.

The method described in III-D can be applied to any of the n images showing the CSs which moved, denoted as \mathbf{M}_i (with $i = 1, \dots, n$), to estimate the position of the floating roof at the corresponding acquisition time. The initial position of the roof could therefore be simply estimated using the first

image (\mathbf{M}_1). However, a more accurate and robust estimation is possible by jointly using all of these n images: because the displacements of the roof are known, the initial position can be obtained from each of these n images by subtracting the displacements accordingly. If we denote as $f_{m_i}(l)$ (with $i = 1, \dots, n$) the functions with the cross-correlation values obtained from each of the n \mathbf{M}_i images, these can be combined to estimate the layover l_{r_1} due to the initial position of the floating roof:

$$l_{r_1} = \arg \max_{l \in [0, l_r^{max}]} \sum_{i=1}^n f_{m_i}(l - d'_{1i}) \quad (23)$$

Here, d'_{1i} denotes the vertical displacement of the floating roof between the first image and image i , which can be computed as:

$$d'_{1i} = \begin{cases} 0 & \text{if } i = 1 \\ \sum_{k=1}^{i-1} d'_{1k} & \text{if } i > 1 \end{cases} \quad (24)$$

When evaluating $f_{m_i}(l - d'_{1i})$ in (23), it is important to note that the functions $f_{m_i}(l)$ have been computed using cross-correlation only inside the valid interval $l \in [0, l_r^{max}]$, and will have a value of 0 outside of this interval.

Finally, the amount of oil stored inside the tank can then be computed for each of the images in the series. If the floating roof did not move, it will have a constant layover value l_{r_s} . On the other hand, if it moved its layover for any given image i in the series can be expressed as: $l_{r_i} = l_{r_1} + d'_{1i}$. In both cases, these layover values must be converted into the corresponding heights, and the volume of oil stored can then be computed using (9).

E. Classification of storage tank type

In section III-E a method was introduced to automatically determine whether a given storage tank has a floating or a fixed roof. This method uses a machine learning classifier trained with three features obtained when estimating the dimensions of a storage tank from a single SAR image. These three features represent the number of CSs which match the two detected semicircular double reflections towards near-range (present in both types of tanks), and the one towards far-range (present only in tanks with a floating roof).

When using a SAR time series the classification approach will be similar, but the available temporal information can now also be used to achieve a better classification performance. If a floating roof moves at some point during the time series, the corresponding CSs will move as well, whereas on a fixed roof all or most of the CSs will remain static. Therefore, the value of w_{avg} (defined in section IV-D) will allow to very easily identify tanks with a floating roof if their roofs move at some point during the time series. If a floating roof did not move, the value of w_{avg} will be low, but the value of $f_s(l_{r_s})$ (also defined in section IV-D) can be used instead. The maximum of these two values can therefore be used as a feature for the classifier, for which a high value will indicate the presence of a floating roof, and a low value a fixed roof. Finally, similarly to the single image case, the number of static CSs matching the double reflections at the bottom and the top of a storage

tank, given by $h(x_b, y_b, r_t)$ and $h(x_b, y_b - l_t, r_t)$, will also be used as the two final features. For these two last features, a small difference with the single image case is that now the Hough accumulator $h(x, y, r)$ was computed using only the CSs which remain static during the series. In summary, when estimating oil storage using a time series, the following three dimensional feature vector \mathbf{v} can be generated for each storage tank:

$$\mathbf{v} = (\max(w_{avg}, f_s(l_{r_s})), h(x_b, y_b, r_t), h(x_b, y_b - l_t, r_t)) \quad (25)$$

As in the single image case, the use of a small number of features will imply that just a few training samples suffice to train an accurate classifier. In section V-C, the performance of the classification using a time series will be evaluated and compared to the case in which a single image is available.

V. RESULTS

To evaluate the performance of the methods presented in this paper, a dataset consisting of three TerraSAR-X repeat-pass images of the port of Rotterdam will be used. These three images have been acquired on July 23, August 3 and August 14, 2017; using the Staring Spotlight imaging mode (with a resolution of 58 cm in slant-range and 23 cm in azimuth) and with an incidence angle of 48.1°. The imaged scene contains 167 oil storage tanks of different sizes, among which there are tanks with floating and fixed roofs. The relevant parameters of each of these storage tanks were obtained by applying the two methods introduced in sections III and IV, using the first image and the complete time series as inputs, respectively. Unfortunately, no ground truth data regarding the sizes of these storage tanks or the amount of oil stored in them was available. However, the same measurement principle applied here (using the semicircular double reflections as described in section II) has also been applied and validated in [15], where the authors manually selected several image points on these double reflections and compared the dimensions derived from these points with those provided by the company operating the oil tanks. Therefore, if the semicircular double reflections are accurately detected by the method proposed in this paper, it will imply that the estimated dimensions should be accurate.

As explained in section III-B, the initial detection of these storage tanks in the imaged scene is out of the scope of this paper, as several works dealing with this topic have already been published. The approximate locations of these tanks were instead automatically obtained using OpenStreetMap data. In addition to their locations, the approximate radius of each tank, which is not needed by the proposed methods but can be used if available, was also extracted from the OpenStreetMap vector data. Then, the resulting GIS data was cleaned by automatically removing those tanks which were either partially outside of the bounds of the imaged scene or too small to be considered of interest. In this case, the few tanks with a radius smaller than 10 m were considered too small and removed from the data, as their size is negligible compared to the many larger tanks with radii up to 50 m, and at those very small sizes the proposed method is more prone to making errors. Finally, in order to evaluate the performance of the proposed method

TABLE I
RESULTS OF THE VISUAL ACCURACY ASSESSMENT

Method	Radius	Roof type	Correct	Wrong	Percent correct
Image	10 – 50	Floating	88	8	91.66%
		Fixed	54	17	76.05%
Image	$\hat{r}_t \pm 5$	Floating	90	6	93.75%
		Fixed	62	9	87.32%
Series	10 – 50	Floating	91	5	94.79%
		Fixed	61	10	85.91%
Series	$\hat{r}_t \pm 5$	Floating	93	3	96.87%
		Fixed	65	6	91.54%

free from the influence of errors in the OpenStreetMap data, the OpenStreetMap data was verified by visually comparing it to the SAR images in a GIS software, and a couple of storage tanks which were present in OpenStreetMap but not in the used SAR images were manually removed.

In the following parts of this section, the accuracy of the obtained results will be analyzed. First, the estimation of the dimensions of the oil storage tanks will be evaluated both qualitatively (by performing a visual accuracy assessment) and quantitatively (by comparing the obtained values with those obtained in manual measurements). Then, the performance of the proposed classification approach to distinguish between tanks with a fixed and a floating roof will be evaluated and compared for different classifiers and different amounts of training data. Finally, a short overview of the method's runtimes will be provided.

A. Visual accuracy assessment

Initially, due to the large amount of oil storage tanks in the scene and the lack of ground truth data, visual accuracy assessment will be performed. For this, the detected semicircular double reflections are drawn over the SAR amplitude image. Only half of each detected semicircle will be drawn, as this makes it easier to determine whether it accurately matches the actual semicircular double reflection on the SAR image. The resulting images are then analyzed by visual inspection to determine the number of tanks for which all the double reflections have been correctly detected, as this will imply that all the parameters have been correctly estimated. The results of this visual accuracy assessment are summarized in Table I.

In this table, results are shown separately for tanks with a fixed and a floating roof, in order to assess whether the method performs better for a specific type. Additionally, results obtained using a single image and the complete time series are compared. Both of these methods were applied twice: once without any prior information about the radius of each storage tank (i.e., radius can have any value between 10 and 50 m), and once using the approximate radius value \hat{r}_t obtained from OpenStreetMap to limit the possible radius values to a smaller interval given by $\hat{r}_t \pm 5$ m.

These results show that all the variants of the proposed method perform very well for storage tanks with a floating roof (which, as previously mentioned, represent the most relevant




Fig. 11. Visualization of the results for many oil tanks of different sizes for a single date, demonstrating the method's performance. Only half of each semicircle is drawn to make the comparison to the SAR amplitude image easier. The ten highlighted tanks will be used for a quantitative accuracy analysis.

use case), with all parameters being correctly estimated for 91.66% to 96.87% of the tanks depending on the method variant. For tanks with a fixed roof the percentage of tanks for which all parameters are correctly estimated is lower, especially for some of the method variants, with this percentage being only 76.05% when using a single image and no prior information about the tank radius, but increasing up to 91.54% when using a time series and the approximate radius value from OpenStreetMap. A possible explanation for the lower performance with tanks with fixed roofs is that their SAR signatures have fewer CSs which can be used for the detection of the semicircular double reflections. Additionally, for this particular scene, the tanks with a fixed roof are mostly of smaller sizes, which further reduces the number of CSs detected in them.

It can also be clearly seen how the use of a time series instead of a single image consistently improves the results for both types of tanks, increasing the number of tanks for which all parameters are correctly estimated. The biggest improvement is obtained for tanks with a fixed roof and when the approximate radius of each tank is a priori unknown, which is the most challenging case.

Finally, Table I also shows how, if available, using an approximate value for the radius of each tank also improves the results, especially for tanks with a fixed roof. This prior knowledge about the tank radii eliminates most of the errors caused by the custom Hough transform detecting semicircles of the wrong radius.

The results obtained using the time series and the approximate radius value from OpenStreetMap can be visualized in

Fig. 11 for a subset of the imaged scene which contains over 50 storage tanks. In it, it can be seen how the semicircular double reflections have been correctly detected for all but one of the storage tanks. This tank has a floating roof and is located towards the bottom left of the image, and the double reflection at its bottom (shown in blue) was wrongly detected at a false position. This implies that the maximum capacity and the amount of oil stored in that particular tank will be wrongly estimated. However, the changes in the amount of oil storage will be correctly estimated anyway, as the radius and the vertical positions of the floating roof were correctly estimated. While Fig. 11 only shows the results obtained for the first image in the time series, the vertical location of the floating roofs was correctly detected for all the storage tanks in all three images. This illustrates an additional advantage of using a time series: the relative changes in the amount of oil stored in tanks with a floating roof will almost always be correctly estimated, even if there are errors when estimating the exact amount of oil or their maximum capacity. This is due to the fact that the roof displacements are estimated in a very robust way, and independently from the detection of the semicircular double reflection of the floating roof.

B. Quantitative accuracy analysis

After the qualitative evaluation of the obtained results via visual analysis, a more detailed analysis of the accuracy achieved will be performed for 10 of the storage tanks shown in Fig. 11 (those enclosed by white rectangles and labelled with the corresponding numbers). Tanks 1 to 5 have a floating

TABLE II
RESULTS FOR THE ESTIMATION OF THE TANK SIZE

Tank	Single Image		Time Series		Manual measurement	
	r_t	h_t	r_t	h_t	r_t	h_t
1	39.71	21.11	40.96	21.11	40.73	21.79
2	44.33	19.07	44.33	19.07	45.36	18.84
3	37.45	21.79	36.95	21.79	38.03	22.13
4	25.32	22.48	24.83	22.48	25.16	23.04
5	38.28	21.79	37.78	21.79	38.26	22.02
6	29.20	14.98	28.96	14.98	27.33	15.32
7	25.23	16.35	25.23	16.35	25.25	17.03
8	20.44	19.07	21.19	18.39	21.23	18.73
9	27.48	14.98	27.73	15.66	25.36	17.93
10	25.48	12.94	25.48	12.94	25.26	12.26

roof, whereas tanks 6 to 10 have a fixed roof, and all of them have different sizes. As no ground truth data is available regarding the dimensions of these storage tanks or the amount of oil stored in them, this accuracy analysis will be performed by comparing the dimensions automatically estimated by the proposed methods (using a single image and a time series) with those obtained from manual measurements. These manual measurements will be performed as described in [15], as this measurement principle has been validated. For this comparison the use of approximate radius information from OpenStreetMap will not play any role, as it only reduces the number of tanks for which the method does not work properly by setting a tighter limit for the possible values of a tank's radius, and it does not have any effect in the accuracy of the estimated values for the cases in which the method works properly.

Table II shows the comparison for the tank radius r_t and height h_t (both given in meters). Here, small differences can be seen between the results obtained from a single image and those obtained from the time series, but it is unclear whether one of these methods is more accurate than the other. The reason for this is that both methods apply a Hough transform in a very similar way for estimating these two parameters. Additionally, the results obtained from these two methods differ slightly from those obtained from manual measurements, which are performed simply by selecting a couple of image points, and therefore they are not necessarily more accurate. However, except for some small differences, all the estimations from the different methods seem to agree. To put these differences into perspective, an error of one pixel will translate into an error of 0.68 m in the estimation of h_t , and of 0.32 m (which is the corresponding square pixel size on the ground) in r_t . Therefore, most of these differences between the different methods are in the order of one to a couple of pixels, with these differences being often slightly larger for the estimated radius than for the height.

In addition to the radius and height of each tank, for the tanks 1 to 5 the vertical position of its floating roof at the time of each image acquisition must be estimated. When applying the method to a single image (the first one in the series) only the initial height of the floating roof h_{r_1} was estimated. When applying the method to the entire time series,

TABLE III
RESULTS FOR THE ESTIMATION OF THE FLOATING ROOF HEIGHT

Tank	Single Image	Time Series			Manual Measurement		
	h_{r_1}	h_{r_1}	h_{r_2}	h_{r_3}	h_{r_1}	h_{r_2}	h_{r_3}
1	4.09	6.81	6.81	6.81	6.83	6.83	6.83
2	8.17	8.17	17.70	15.66	10.09	19.63	17.47
3	12.26	10.90	10.90	23.16	13.31	13.31	25.57
4	6.13	5.45	14.98	14.98	6.51	16.16	16.16
5	17.03	17.71	17.71	21.11	18.91	18.91	22.55

and in the manual measurements performed for comparison, the height of the floating roof was estimated for all three images. The obtained results can be seen in Table III. Again, the results obtained from the manual measurements are not necessarily more accurate, but they serve as useful reference for this comparison in the absence of ground truth data. In this table, some significant differences can be observed across the different estimates for the initial height of the floating roofs. On the other hand, the vertical displacements of the floating roof in between consecutive images (i.e., the difference between $h_{r_{i+1}}$ and h_{r_i}) estimated with the proposed method agree very well with the manual measurements. The main cause of these differences in the initial height is actually the difference in the previously estimated radius r_t (shown in Table II). For the used incidence angle, an error of 1 m in r_t will induce an error of 2.23 m in h_{r_1} . This is caused by the fact that the semicircular reflection of the floating roof appears towards far-range, whereas the two reflections at the outer tank structure appear toward near-range. Therefore, this is not a particular issue of the proposed method, but it will be common to all the methods employing SAR images.

In conclusion, the estimates for r_t , h_t and h_{r_1} appear to agree between the proposed method (both using a single image and a time series) and the manual measurements, but with small differences in the range of one to a couple of pixels. On the other hand, the vertical displacements of the floating roof between consecutive images seem to be estimated much more accurately. This implies that the relative changes in the amount of oil stored in a tank with a floating roof will be estimated very accurately, whereas absolute measurements of the amount of oil stored in it will be slightly less accurate. This also applies to the estimated maximum capacity, both for tanks with fixed and floating roofs. The accuracy of these absolute estimates could potentially be improved by exploiting the fact that storage tanks often have standard sizes, rounding the estimates for r_t and h_t to the closest values corresponding to a standard size. This would in turn also remove the errors induced in the estimation of h_{r_1} due to an inaccurate radius.

C. Classification performance

To evaluate the performance of the proposed classification approach to distinguish between tanks with a fixed and a floating roof, a dataset had to be first generated by assigning the appropriate label to each of the storage tanks, as this information was not available in the used OpenStreetMap data. These labels were manually assigned by visualizing each

storage tank in the SAR images, resulting in a dataset with 167 samples which will then be split into training and test data.

For each of the storage tanks in this dataset the corresponding features were computed from both a single image and the time series, as described in sections III-E and IV-E, respectively. The classification results using these two different sets of features (obtained from a single image and a time series) were compared for different amounts of training data and two different classifiers: a support vector machine (SVM) and a random forest. To purely evaluate the classification without the influence of the errors occurred during the previous steps, the few tanks for which the estimation of the tank dimensions did not work properly (i.e., those listed in Table I as wrong) were removed from the dataset.

In order to analyze the effect that the amount of training data has in the classification results, both the SVM and random forest classifiers were trained multiple times from scratch using training datasets of different sizes (with 10, 25, 50, 75 and 100 samples). Each time the training samples were selected randomly while keeping an equal split of tanks with floating and fixed roofs, and the remainder of the dataset was used as test data. To minimize the influence that the selection of the training samples has in the results, this process was repeated 100 times for each case, and the F_1 score of the corresponding 100 trained classifiers was averaged, obtaining an average F_1 score for each number of training samples. The results obtained for both classifiers and both sets of input features (obtained from a single image and from the time series) can be seen in Fig. 12. As expected, due to the use of just three simple but very informative features (both when using a single image as well as the time series) the two classifiers perform very well even with small training datasets, obtaining average F_1 scores above 0.9 with just 10 training samples. The best classification results are consistently obtained for all the different amounts of training data when using the input features obtained from the time series and a random forest classifier, which reaches an average F_1 score of 0.97 with 100 training samples. When using the input features obtained from a single image, the SVM classifier performs slightly better than the random forest, reaching an average F_1 score of 0.95 with 100 training samples.

D. Runtime analysis

The proposed method could potentially be applied to monitor the amount of oil stored in many oil refineries, each with a large number of storage tanks, regularly every few days. Here a brief analysis of the runtime of the proposed method will be provided, to give an outlook on how the processing time scales up with the number of SAR images and/or number of storage tanks to be monitored. This processing time will be analyzed separately for two separate stages of the proposed method: pre-processing (which involves the detection of CSs in the SAR images, as well as change detection when using a time series), and information extraction (which involves the automatic estimation of all the relevant parameters for each storage tank).

For the pre-processing stage the runtime will depend on the number of SAR images and their sizes. When using a




Fig. 12. Comparison of the classification performance using the features extracted from a single image and a time series for two different classifiers and different amounts of training data.

single SAR image the pre-processing involves just the CS detection which, for one of the used TerraSAR-X images with 9014×18847 pixels, takes 2 minutes and 49 seconds on a Intel i7-8665U laptop CPU (with 4 cores) and 32 GB of RAM. When using a time series the pre-processing also involves co-registration and coherent change detection in addition to the CS detection. For the series of three TerraSAR-X images used in this paper, this takes a total of 10 minutes and 48 seconds (or 3 minutes and 36 seconds per image) on that same hardware. In both cases, the pre-processing algorithm was implemented in Python, using Numpy compiled with Intel MKL. In conclusion, the pre-processing stage appears to be fast enough that many SAR images could be processed in a reasonable time, even when running on consumer hardware. If required, faster pre-processing could be achieved by simply using a more powerful CPU or by implementing the algorithms on a GPU.

Regarding the information extraction stage, its runtime will mainly depend on the number of storage tanks in the imaged scene, as the processing will be performed separately for each storage tank using just a small image patch around it. As with the pre-processing stage, the code was implemented in Python using Numpy, and it was tested on the same hardware. When using a single SAR image and the generic interval for the tank radius (10 – 50 m), the runtime was 27 minutes and 18 seconds for the 167 storage tanks in the imaged scene, or 9.80 seconds per storage tank. If the approximate radius \hat{r}_t for each tank (obtained from OpenStreetMap) is used to define a tighter interval for the radius ($\hat{r}_t \pm 5$), the runtime becomes much faster: 2 minutes and 54 seconds in total, or 1.04 seconds per tank. The reason for this large difference in speed is that increasing the interval for the radius increases the computation time of the Hough transform (which is the most time consuming step), and it also makes the size of the image patch surrounding each tank larger, which further increases the runtime. When using a time series, the information extraction stage runs faster than when using a single image, even if multiple images are being processed at once. The reason for this is that the computation of the Hough transform

becomes faster, as only the CSs which remain static are taken into account. For the generic radius interval, the information extraction from the three TerraSAR-X images takes 19 minutes and 50 seconds, or 6.12 seconds per tank. If the approximate radius from OpenStreetMap is used, the processing takes just 2 minutes and 30 seconds, or 0.89 seconds per tank.

It is important to note that when using a time series, the whole processing does not need to be repeated each time that a new image is acquired. For the pre-processing stage, the new image can be simply co-registered and added to the existing image stack, and the change detection results can be updated including the coherence of the new image pair and the CSs detected in this new image. Regarding the information extraction stage, the Hough transform does not need to be computed again, as the location, radius and height of the storage tanks will not change. Instead, only the displacements of the floating roof and its initial position can be updated taking into account the new image.

VI. CONCLUSION

The monitoring of oil storage has become a popular application of high resolution SAR and optical remote sensing data, with multiple methods presented in the literature for the detection of oil storage tanks and the estimation of their relevant dimensions. Unlike the methods employing optical images, the methods using SAR images can ensure regular observations with optimal accuracy, as they are not affected by external factors (e.g., sun and clouds). In this paper, a method has been introduced for the monitoring of oil storage tanks using one or more high resolution SAR images. The proposed method has significant advantages with respect to the state-of-the-art, mainly: it works with both tanks with fixed and floating roofs and can automatically distinguish them, it can fully automatically extract all their relevant parameters by exploiting their complete geometry (rather than just using a few select image points), and if a time series is available all the observations can be processed jointly to obtain more robust and accurate estimates of all the parameters (and especially of the vertical displacements of floating roofs). This method has been applied to three TerraSAR-X images of the port of Rotterdam containing 167 storage tanks of different sizes and types.

Before all the relevant parameters of a storage tank can be estimated, its approximate location in the image should be known. This information was obtained from OpenStreetMap data, as the pre-detection of the storage tanks is outside of the scope of this paper. Alternatively, this information could also be extracted from the available imagery using any of the methods available in the literature, many of which have been listed in section III-B.

An analysis of the obtained results has shown that as expected, when the proposed method is applied to a time series, the exploitation of the available temporal information increases the robustness of the method, reducing the number of tanks for which the method provides wrong results, especially for tanks with a fixed roof and/or smaller sizes. Additionally, for tanks with a floating roof, the use of a time series has

another advantage: the vertical displacements of the floating roof (which are of great interest as they give the relative changes in the amount of oil stored) can be estimated much more accurately (even with subpixel accuracy) and robustly than when multiple images are processed individually.

The results have also shown that if available, the use of some information regarding the approximate radius of each tank is of advantage, as this reduces the size of the interval for the possible radius values, which leads to fewer errors and also greatly improves the runtime. In our case these approximate radius values were obtained from OpenStreetMap, but these could also be obtained from the available imagery by some of the methods listed for the pre-detection step. Alternatively, rather than trying the many possible radius values covering a given interval, only a few specific values for the radius could be used, as most storage tanks have standard sizes.

When using a time series and the approximate radius information, the method performed very well, correctly fitting all the semicircular double reflections for 158 of the 167 storage tanks in the imaged scene, with errors in 3 of the 96 tanks with a floating roof, and 6 of the 71 tanks with a fixed roof. Overall, the method appears to perform better for tanks with a floating roof, which represent the most interesting use case, as the amount of oil stored in a given tank can only be determined if it has a floating roof.

A good performance in the classification of storage tanks (i.e., whether they have a fixed or a floating roof) has also been demonstrated, achieving average F_1 scores above 0.9 with as little as 10 training samples. Here, the use of a time series also slightly improves the classification results, especially with a larger training dataset, achieving an average F_1 score of 0.97 with 100 training samples. An analysis of the method's runtime has also been performed, showing that the method runs fast enough even on a standard laptop, making the processing of a large amount of SAR images containing lots of storage tanks feasible.

Further work should involve a more rigorous verification of the accuracy achievable by the proposed method. For this ground truth data should be obtained for several storage tanks (including their dimensions and the amount of oil stored in them at different times), and SAR images of the same scene should be acquired at these same times. If during this validation a bias were to be found in some of the estimates (e.g., such as the tank height being very likely slightly underestimated as previously mentioned in section II-A), the ground truth data could also potentially allow to derive some correction factors to compensate for this bias. Additionally, the potential of combining SAR images acquired with different imaging geometries (e.g., different orbits and/or incidence angles) could also be investigated. While the proposed method can already be applied separately to multiple images or time series acquired with different geometries to achieve more frequent observations, two such images could also be exploited jointly to improve the accuracy of the estimates for each tank's radius and height. Such improved estimates could be achieved by fitting the two semicircular double reflections of the outer tank structure jointly using both images, adding in this way an additional geometric constraint.

ACKNOWLEDGMENT

The authors would like to thank M. Hager and H. Anglberger for their help with the acquisition of the TerraSAR-X data used in this paper, and the anonymous reviewers for their constructive comments and recommendations. The TerraSAR-X images were provided in the frame of the science proposal MTH3256.

REFERENCES

- [1] S. Pullarcot, *Above Ground Storage Tanks: Practical Guide to Construction, Inspection, and Testing*. CRC Press, 2015.
- [2] *API Standard 650: Welded Tanks for Oil Storage*, 12th ed., American Petroleum Institute, March 2013.
- [3] "Orbital Insight Oil Inventories." [Online]. Available: <https://orbitalinsight.com/use-cases/oil-inventories> (last accessed on March 9th, 2021).
- [4] "Ursa Space Global Oil Storage." [Online]. Available: <https://www.ursaspace.com/solutions/direct-measurements> (last accessed on March 9th, 2021).
- [5] A. O. Ok and E. Baseski, "Circular Oil Tank Detection From Panchromatic Satellite Images: A New Automated Approach," *IEEE Geoscience and Remote Sensing Letters*, vol. 12, no. 6, pp. 1347–1351, 2015.
- [6] L. Zhang, Z. Shi, and J. Wu, "A Hierarchical Oil Tank Detector With Deep Surrounding Features for High-Resolution Optical Satellite Imagery," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 10, pp. 4895–4909, 2015.
- [7] Z. Liu, D. Zhao, Z. Shi, and Z. Jiang, "Unsupervised Saliency Model with Color Markov Chain for Oil Tank Detection," *Remote Sensing*, vol. 11, no. 9, p. 1089, 2019.
- [8] M. Jing, D. Zhao, M. Zhou, Y. Gao, Z. Jiang, and Z. Shi, "Unsupervised Oil Tank Detection by Shape-Guide Saliency Model," *IEEE Geoscience and Remote Sensing Letters*, vol. 16, no. 3, pp. 477–481, 2019.
- [9] Q. Wang, J. Zhang, X. Hu, and Y. Wang, "Automatic Detection and Classification of Oil Tanks in Optical Satellite Images Based on Convolutional Neural Network," in *Image and Signal Processing*, ser. Lecture Notes in Computer Science, A. Mansouri, F. Nouboud, A. Chalifour, D. Mammas, J. Meunier, and A. Elmoataz, Eds. Cham: Springer International Publishing, 2016, vol. 9680, pp. 304–313.
- [10] Q. Wang, J. Zhang, and X. Hu, "Automatic Oil Reserve Analysis Through the Shadows of Exterior Floating Crest Oil Tanks in Highlight Optical Satellite Images," in *Advances in Visual Computing*, ser. Lecture Notes in Computer Science, G. Bebis, R. Boyle, B. Parvin, D. Koracin, F. Porikli, S. Skaff, A. Entezari, J. Min, D. Iwai, A. Sadagic, C. Scheidegger, and T. Isenberg, Eds. Cham: Springer International Publishing, 2016, vol. 10073, pp. 23–32.
- [11] H. Xu, W. Chen, B. Sun, Y. Chen, and C. Li, "Oil tank detection in synthetic aperture radar images based on quasi-circular shadow and highlighting arcs," *Journal of Applied Remote Sensing*, vol. 8, no. 1, p. 083689, 2014.
- [12] L. Zhang, S. Wang, C. Liu, and Y. Wang, "Saliency-Driven Oil Tank Detection Based on Multidimensional Feature Vector Clustering for SAR Images," *IEEE Geoscience and Remote Sensing Letters*, vol. 16, no. 4, pp. 653–657, 2019.
- [13] L. Zhang and C. Liu, "Oil Tank Extraction Based on Joint-Spatial Saliency Analysis for Multiple SAR Images," *IEEE Geoscience and Remote Sensing Letters*, vol. 17, no. 6, pp. 998–1002, 2020.
- [14] R. Guida, A. Iodice, and D. Riccio, "Assessment of TerraSAR-X Products with a New Feature Extraction Application: Monitoring of Cylindrical Tanks," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 48, no. 2, pp. 930–938, 2010.
- [15] H. Hammer, S. Kuny, and K. Schulz, "Simulation-Based Signature Analysis of Fuel Storage Tanks in High-Resolution SAR Images," *IEEE Geoscience and Remote Sensing Letters*, vol. 14, no. 8, pp. 1278–1282, 2017.
- [16] T. Anahara and M. Shimada, "Oil Storage Estimation with Time-Series L-Band Sar Imagery," in *2018 IEEE International Geoscience & Remote Sensing Symposium*. Piscataway, NJ: IEEE, 2018, pp. 842–845.
- [17] M. Back and T. Jeon, "Analysis of Oil Storage Trend Using KOMPSAT-5 SAR Data," in *IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 9262020, pp. 4088–4091.
- [18] M. J. Sanjuan-Ferrer, I. Hajnsek, K. P. Papathanassiou, and A. Moreira, "A New Detection Algorithm for Coherent Scatterers in SAR Data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 11, pp. 6293–6307, 2015.
- [19] R. Z. Schneider, K. P. Papathanassiou, I. Hajnsek, and A. Moreira, "Polarimetric and interferometric characterization of coherent scatterers in urban areas," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 44, no. 4, pp. 971–984, 2006.
- [20] R. Z. Schneider and K. Papathanassiou, "Estimation and correction of ionospheric induced phase errors in SAR images using Coherent Scatterers," in *2009 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 12.07.2009 - 17.07.2009, pp. IV–165–IV–168.
- [21] V. M. Giacomazzo, A. Refice, F. Bovenga, and N. Veneziani, "Identification of Coherent Scatterers: Spectral Correlation vs. Multi-Chromatic Phase Analysis," in *IGARSS 2008 - 2008 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 07.07.2008 - 11.07.2008, pp. IV - 411–IV - 414.
- [22] W. Guo, W. Yang, H. Zhang, and G. Hua, "Geospatial Object Detection in High Resolution Satellite Images Based on Multi-Scale Convolutional Neural Network," *Remote Sensing*, vol. 10, no. 1, p. 131, 2018.
- [23] A. Ferretti, C. Prati, and F. Rocca, "Permanent scatterers in SAR interferometry," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 1, pp. 8–20, 2001.
- [24] M. Eineder, N. Adam, R. Bamler, N. Yague-Martinez, and H. Breit, "Spaceborne Spotlight SAR Interferometry With TerraSAR-X," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 5, pp. 1524–1535, 2009.
- [25] E. Sansosti, P. Berardino, M. Manunta, F. Serafino, and G. Fornaro, "Geometrical SAR image registration," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 44, no. 10, pp. 2861–2870, 2006.



Carlos Villamil Lopez was born in Cervo, Spain, in 1989. He received the M.Sc. degree in telecommunications engineering from the University of Oviedo, Gijón, Spain, in 2013. He is currently working towards the Ph.D. degree at the Department of Photogrammetry and Remote Sensing, Technical University of Munich (TUM), Munich, Germany.

In 2012, he was a Visiting Student with Virginia Tech, Blacksburg, VA, USA. Since 2013, he has been with the Microwaves and Radar Institute, German Aerospace Center (DLR), Oberpfaffenhofen, Germany, as a research scientist. His main research interests include synthetic aperture radar (SAR) image processing, and SAR change detection.



Uwe Stilla (M'04-SM'09) was born in Cologne, Germany, in 1957. He received the Diploma (Dipl.-Ing.) degree in communications engineering from University of Paderborn, Germany, the Diploma (Dipl.-Ing.) in electrical engineering and the Ph.D. (Doctor of Engineering) degree in electrical engineering (pattern recognition) from the University of Karlsruhe, Germany, in 1980, 1987, and 1993, respectively. From 1987 to 1990, he was with the Institute of Biomedical Engineering, University of Karlsruhe, and from 1990 to 2004, with the Institute of Optronics and Pattern Recognition (FGAN-FOM, now Fraunhofer IOSB, Ettlingen). Since 2004 Uwe Stilla is Professor at Technical University of Munich (TUM), Germany. He is head of the Unit "Photogrammetry and Remote Sensing" at the TUM Department of Aerospace and Geodesy and was the Dean of Studies (2005-2016) and additionally Vice Dean (2008-2013) of the TUM Department of Civil Engineering and Geodesy. From 2012 to 2016 he was Vice President and from 2016 to 2020 President of the German Society for Photogrammetry, Remote Sensing and Geoinformation (DGPF). Since 2004 he holds the chair of multiple working groups of the International Society of Photogrammetry and Remote Sensing (ISPRS), last: Inter Commission Working Group ICWG II/III 'Pattern Analysis in Remote Sensing'. He is member of the Scientific Board of German Commission of Geodesy (DGK), and Vice Chairman of Commission for Geodesy and Glaciology (KEG) of the Bavarian Academy of Science and Humanities (BADW). His research focuses on image analysis in the field of photogrammetry and remote sensing. He organized and co-organized many conferences and workshops (ISPRS, IEEE, IAPR) and published more than 500 scientific contributions.