299927

2014 INTERNATIONAL OIL SPILL CONFERENCE

A Model Based Approach to Radar Oil Spill Detection

Torstein Pedersen
Nortek AS
Vangkroken 2
1351 Rud, Norway

Javier Perez
Nortek BV
Schipholweg 333a
1171PL Badhoevedorp
Nederland

Jos Van Heseen
Nortek BV
Schipholweg 333a
1171PL Badhoevedorp
Nederland

ABSTRACT 299927:

A typical oil spill recovery vessel has been historically outfitted with an oil spill detection (OSD) radar. During an oil spill recovery operation, there is a dedicated operator who is responsible for interpreting information from the radar image. Industry developments over the last several years now require that an OSD radar *automatically* detect and track an oil spill. There are two primary needs driving this development. The first is that OSD systems and operations are becoming more sophisticated; automatic OSD aids for a more efficient oil spill operation where an operator's attention may be directed to a potential spill. The automatic OSD also aids a multi-sensor system; one such example is where an OSD radar is used to steer an IR camera to a candidate spill for more detailed evaluation or validation. The other primary driver for automatic OSD is for monitoring systems, which serve for early warning. Monitoring systems may be found along coastal installations or oil platforms.

The automatic spill detection functionality of an OSD system may be implemented in different levels of sophistication. Perhaps the simplest configuration is one that uses fixed thresholds relative to the image for alarming whether a region in a radar image is a spill or not. The benefit of simple threshold detector is that it is easy to implement in software. The weakness is that it is prone to both lower overall detection rate and high false alarm rate. A more robust automatic spill detection method is one that treats it as an image-processing problem. The paper here presents a model based OSD.

Generation of confidence maps is central to the method and provides an indication of the likelihood of oil. Inputs to the confidence maps come from multiple sources, several of which are based on uniquely constructed models. Among these is a histogram comparator, which scans a radar image and compares the data to reference models from real oil spills.

A discussion of the methods used focuses on (a) the necessary steps prior to the confidence map construction, (b) how the confidence maps are layered with inputs, (c) how the

information in the confidence maps is transitioned into the detection of oil, (d) and finally alarming.

INTRODUCTION:

Oil spill detection by way of radars is a proven and tried technology. Vessels dedicated to oil spill recovery have been employing this tool for over 10 years, and its value is most evident during challenging environments or poor visibility. The most elemental description of radar OSD is its ability to exploit the sea clutter and identify regions where the sea clutter is absent, which is the effect of oil on the sea surface. The system is characterized as a dedicated OSD system if it can remove the background signature of the ocean waves and emphasizes the disparity between background sea clutter and oil covered water in the radar image. Like all remote sensing options available today, it is not without its challenges or limitations.

Much has been done to improve upon the limitations whether it has been to extend the detection range or to improve performance during wind calm periods. The radar system's sensitivity has been one area in which we have witnessed specific improvements. The hardware handling the signal is typically low noise analog to digital converters with high digital resolution (upwards to 14 bit analog digital converter-ADC). Another system element that has seen special attention is the front end (antenna); for OSD vertically polarized (VV) antennae have the desired backscatter response – which is counter the tradition of using horizontally polarized (HH) antennae in the marine community.

More recently, the industry has required automation of the task of OSD. Essentially, certain features in a radar image need to be classified without the watchful eye of an operator. Now this may be done by employing simple thresholds upon the image's greyscale as to what one identifies as oil or not. However the skill of automatic OSD improves with a more advanced and rigorous treatment of the problem. The authors view this as an image-processing problem.

This approach focuses on methods that are not only better at distinguishing oil spills within background clutter, but also to quantifies the *likelihood* of the image feature – such as a spill. This is a model driven approach where clutter characteristics of both water and oil are constructed from a database of real oil spills. Candidate spills are qualified based on the effort required to match the candidate's distribution to the modeled distribution of both oil and water clutter. The process results in a confidence map, which is a pixel-by-pixel indication of the presence of oil likelihood.

DISCUSSION:

Before we delve into a detailed description of the models used for the image processing, it is sensible to establish an understanding of the principles of operation as well as the necessary image preprocessing.

Principle of operation

Radars used for oil spill detection (OSD) exploit the signal that mariners typically try to filter out; this is the *sea clutter* or the backscatter signal from the sea surface. The dominant source of backscatter comes from surface capillary waves, which are generated by winds in the immediate vicinity. The signal wavelength is similar in length to the capillary waves (approximately 3 cm) and the signal constructively reflects from the sea surface. When oil is present on the surface, it suppresses or attenuates the capillary waves because the surface no longer has the surface tension characteristics to support the capillary waves. The result in a radar image is that the oil appears as a region of reduced backscatter, or a dark spot in an otherwise clutter rich image.

Utility of such a system requires local winds to generate capillary waves. Without surface winds, the sea surface will become glassy calm and the radar image will have significantly reduced disparity between clutter and no clutter, or water – oil. A practical figure used in the community is that winds need to be a minimum of 2 m/s and no greater than 12 m/s, which is when wave breaking dominates the sea surface.

Preprocessing

The preprocessing that is required prior to implementing the detection model requires that the image processing techniques have an optimum image to operate. The first step involves correcting for vessel motion for each radar radial line in the image. This is accomplished through the input from the ship's GPS and gyro. The radar data stream is referenced to Earth coordinate, as opposed to ship referenced.

The next task is to remove the ocean wave signature by averaging together consecutive images. This is followed by normalizing the signal intensity with respect to range, which is a compensation for the reduced signal with increased range. Once the image has been cleaned, then the automatic OSD can commence. One should note that these preprocessing steps and the subsequent steps for the automatic detection occur at the image update rate, which is the antenna rotation rate – typically 22, 45, or 60 rpm. Figure 1 provides a more detailed sequence of the preprocessing steps. Figure 2 illustrates the results of selected preprocessing steps; one can clearly see the presence of oil in Figure 1c, which is located due North.

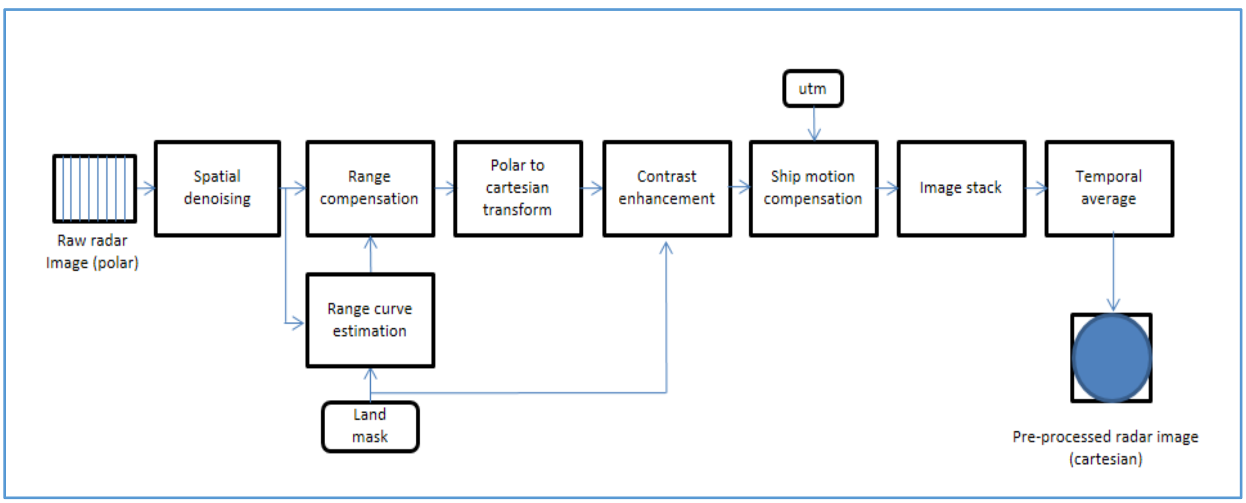

Figure 1 Sequence of preprocessing steps from raw data to final preprocessed image.

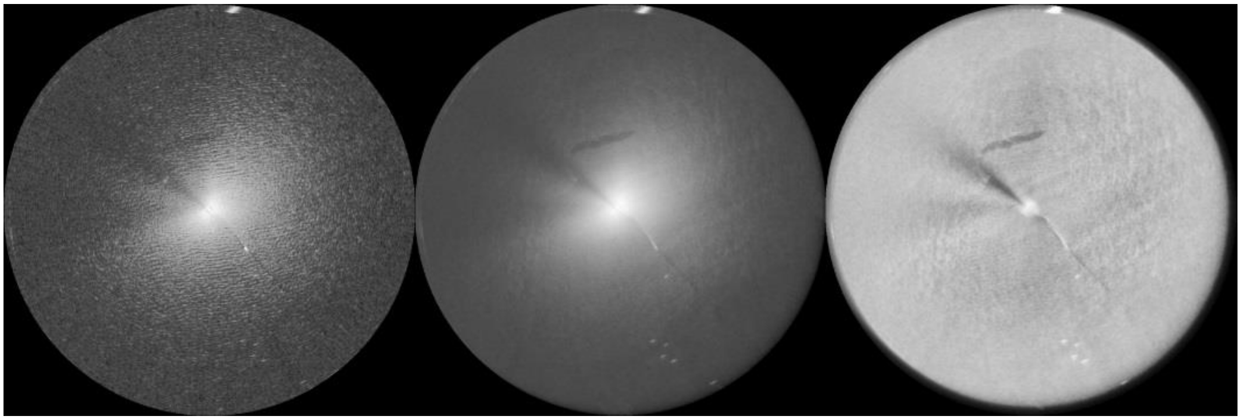

Figure 2 Series of radar images through the steps of preprocessing. Note the improved clarity of the oil spill due North. (a) raw image, (b) ocean waves filtered, (c) after range correction.

The automatic detection of oil can begin now that the image is properly optimized. A simple and often used technique is to apply a simple intensity thresholds: oil may be detected by establishing a lower threshold limit for the pixel intensity; pixel values above the threshold are water and below this surface are oil covered. Once a minimum number of neighboring pixels fulfilling this requirement exist, the region is classified as oil, and the OSD may then automatically alarm that oil is present in the image. The advantage of a simple threshold detector is that it is easy to implement, however it does have its limitations. One shortcoming is that the backscatter intensity values can vary for different environmental conditions (i.e., variable wind strengths) and therefore establishing a fixed threshold can be challenging.

Detection Model

We have three distinct steps that define our process of automatic oil spill detection. These steps naturally follow the preprocessing and are (a) the generation of the confidence maps, (b) spill segmentation, (c) and the validation. This is illustrated in the block diagram of Figure 3, and is further detailed below.

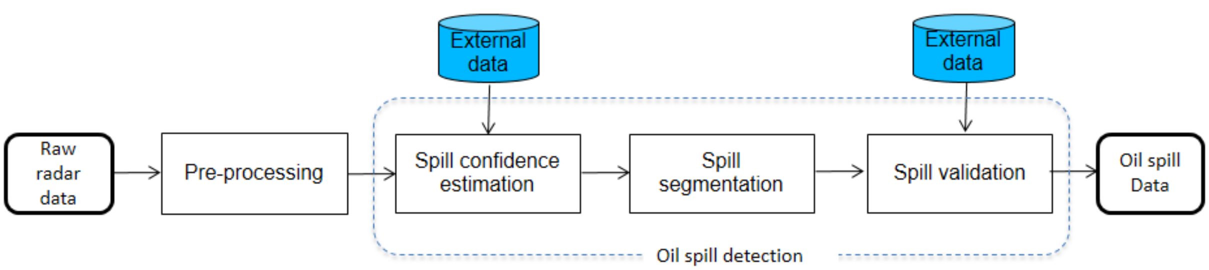

Figure 3 Automatic oil spill detection block diagram.

Confidence Maps

The approach that Nortek uses for oil detection is a model-based approach, and by this, we mean that the problem is treated as an image processing problem. As such the image is treated as one that is composed of characteristic features. This includes sea clutter, oil spills, radar shadows, wind shadows, landmasses, vessels, etc. Comparative processing of these features with models based on real features (e.g. an oil spill) distinguishes the method from the classic threshold setting method.

We begin by looking at the identification of oil within the background sea clutter. An image is scanned over groupings of pixels; for the sake of example let us say this is a 10x10 pixel region. A histogram of the gray scale (intensity) distribution is calculated and this distribution is compared to a histogram of water and a histogram of oil. Examples of histograms for the two sources is presented in Figure 4.

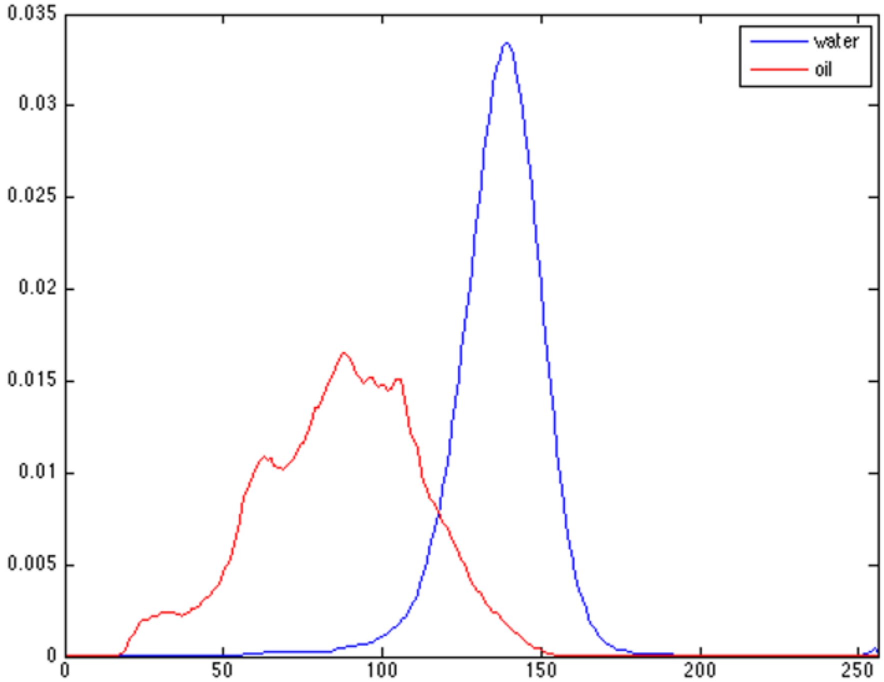

Figure 4 Histograms of water (blue) and oil covered water (red).

Nortek uses a comparison process known as the *Earth Mover Distance* (EMD) method. This is defined as the minimum cost or effort required for two histograms to match. The effort is estimated by the product of the gray scale difference (distance) and the number of observations (amount of "energy"). It is a distance measure between two histograms; the effort is estimated by determining how much inter-bin (of grey scale) distance needs to be exchanged in order to achieve similarity. This is a optimization process where there are several permutations for transferring the grey scale levels. The minimum effort of transferring the different levels over distances is a measure of the similarity to either water or oil. The process accounts for the histogram shape and is not a simple difference of the histograms. Examples of the possible combinations are illustrated in Figure 5. The final estimate is a measure of the oil likelihood and contributes directly to the *Confidence* estimates.

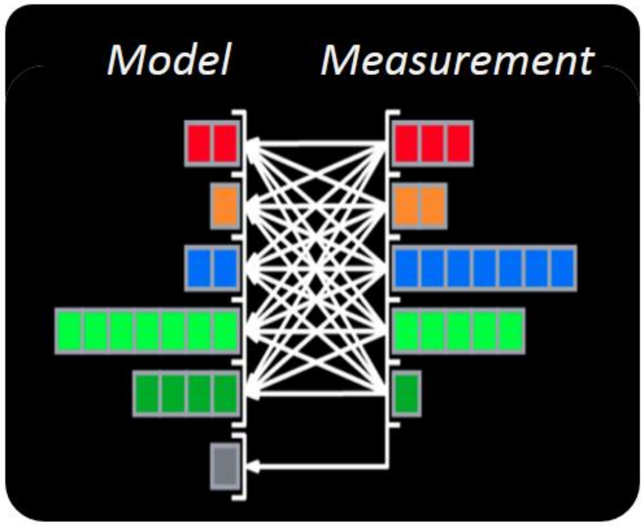

Figure 5 Earth Mover Distance metric illustrating simple histograms comparisons of a model and a measurement.

The challenge with employing the Earth Mover Distance (EMD) method is that it is fairly complex and computationally demanding. This is not so much that a simple histogram comparison is demanding, but that it must be repeated for each pixel which extents for 1024x1024 and the image update can occur as rapid as 1 Hz. It was subsequently deemed that the best way to handle this was to use a Wavelet EMD which is computationally efficient with respect to the achieved accuracy. Figure 6 shows the results of the processing on an image with sea clutter and radar mast shadows. The resulting confidence maps are for processing with floating point EMD, integer EMD, and Wavelet EMD.

The description of the solution so far is the Nortek model-based detector in its simplest form. It effectively handles the problem of OSD in open waters where the field of view is not obscured. This is rarely the case when the oil hits the water. The image may contain several sources of false detections such as land masses or vessels. Minimizing false detects is once again an image processing problem.

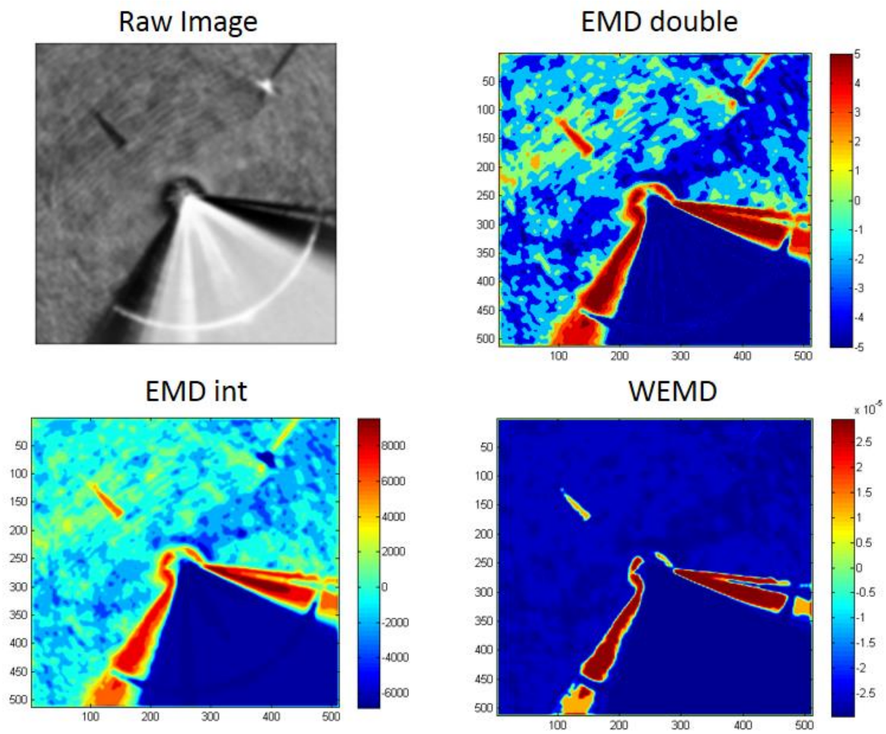

Figure 6 Results of different Earth Mover Distance implementations: double, integar, wavelet.

Landmasses present complications in few different ways. First of all, we know that we are not going to be searching for oil on land, so we neglect land in the EMD. Land quite often has reduced wind velocity along its boundary, and regions of reduced wind could be detected as oil covered water. This is particularly true in the lee of the land but it may also occur on the windward side. In order to handle this the software treats the land as a zero probability for the EMD processing, neglects land from the range compensation, and a buffer region along the coast is set to zero probability for confidence map to avoid wind shadows from detection.

Further to this, land has a large backscatter and complicates the range correction for the pre-processing. The landmass therefore is neglected from the range compensation. The definition of the land boundaries comes from a worldwide land database. An example of the land overlay is presented in Figure 7.

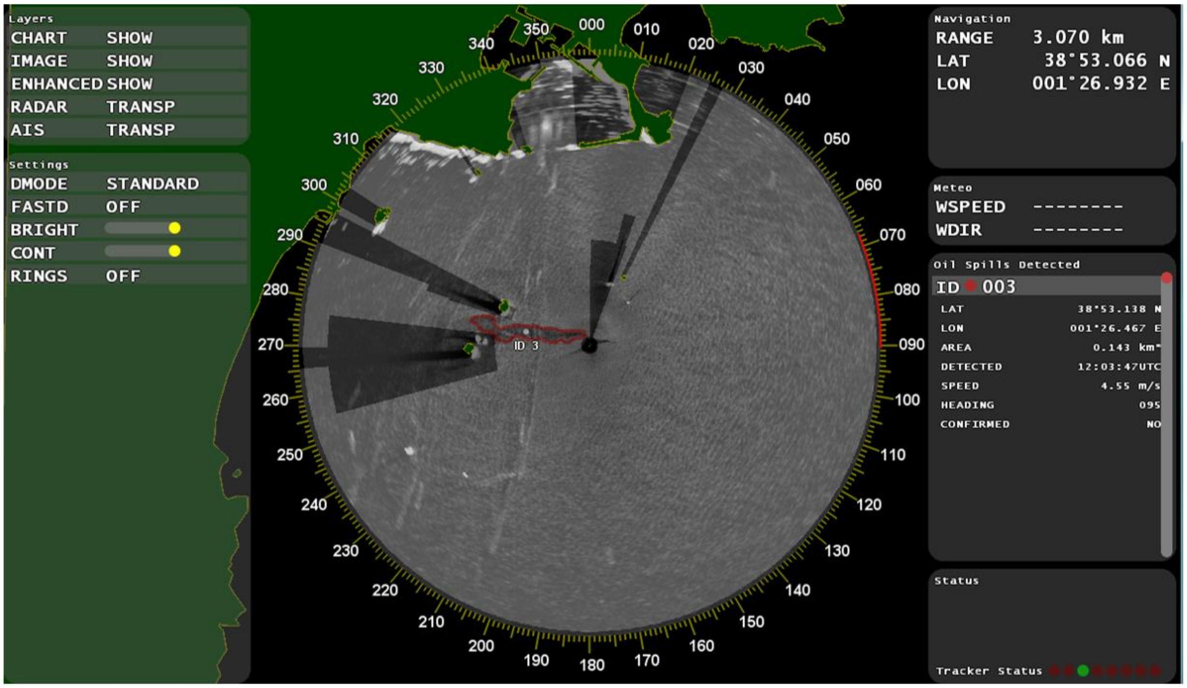

Figure 7 Overlay of the land database in a preprocessed radar image. Note the shadows from various islands are blocked. There is also an oil spill to the west of the center outlined and marked with ID 3.

Radar shadows behind vessels and other obstacles may appear as oil to the EMD processing and therefore these regions are set to zero probability when detected. In the context of constructing a confidence map it is important to understand that this binary probability value is assigned for each of the image frames, and since multiple frames construct the pre-processed image the shadow region becomes a weighted based on the number of images where it was detected. The result from the shadow detector is a weighted confidence.

This requires first correctly identifying the radar shadow. The method Nortek uses is to search for a characteristic feature, which is often a strong return followed by an immediately reduced return further along the radial. Once again, a collection of images containing known vessels with radar shadows was used to define a model for identifying this feature. We note that this is a similar image processing technique that is used for face detection in cameras. A variety of different vessel shadows is presented in Figure 8, and the results of weighted detector is presented in Figure 9.

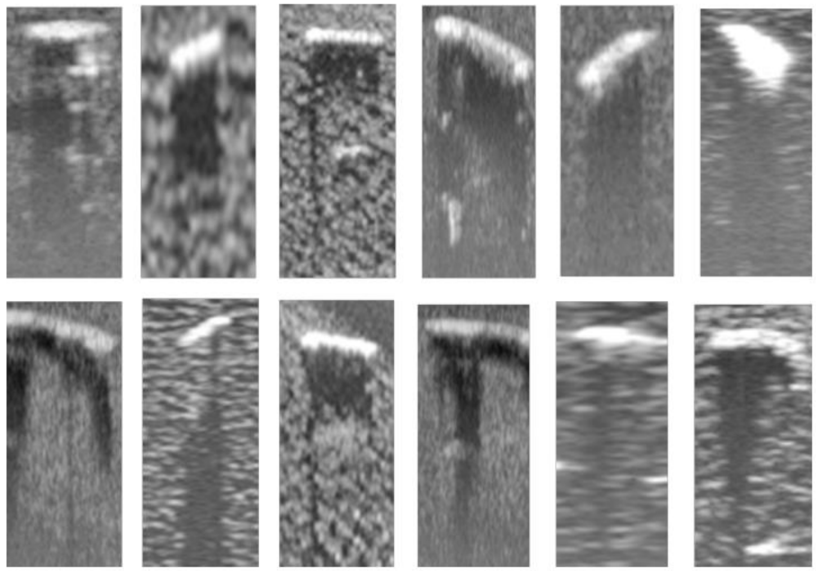

Figure 8 Examples of different vessels and radar shadows used for algorithm training.

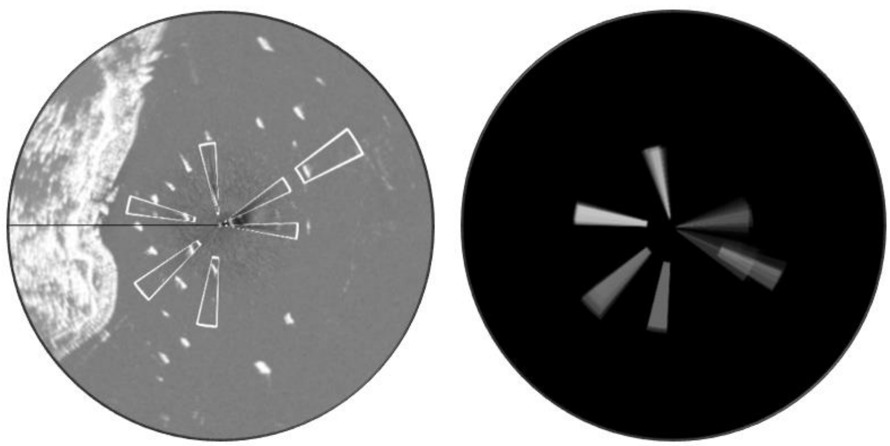

Figure 9 Example of vessel detection and weighting from different radar images.

The vessel detection employs a boosted classifier for the automatic shadow detection. In very simple terms, the complete classifier is strong (robust) classifier which is constructed by a weighted, linear combination of simpler classifier outputs. A simple classifier is not sufficient on its own but is easy to compute. The algorithm (AdaBoost) used here interactively selects a subset of classification functions such that the classification error (i.e., false detect) of the resulting strong classifier is minimized.

Another way to describe this is that the algorithm effectively chains together simple classifiers, where the first stage has a high detection rate (DR) but also a high false acceptance (50%). Successive stages are more refined and improve rejection of the false acceptance with a mild loss of detection rate. At processing time there is a small window scanning through the image. The detection chain reduces complexity but does not compromise skill.

The detection scheme employs a cascade of strong classifiers where the detection rate is designed to be very high (DR=0.99) and the false detection rate (FDR) is allowed also to be high (FDR < 0.5). If one uses many steps in the cascade then the final DR and FDR will both be reduced according to the following.

$$DR = \prod_{j=1}^{N} d_j \qquad FDR = \prod_{j=1}^{N} f_j$$

A cascade of 15 stages can achieve a detection rate of $DR = 0.99^{15} = 0.86$, and the false detection rate as $FDR = 0.5^{15} = 3.0 \times 10^{-5}$.

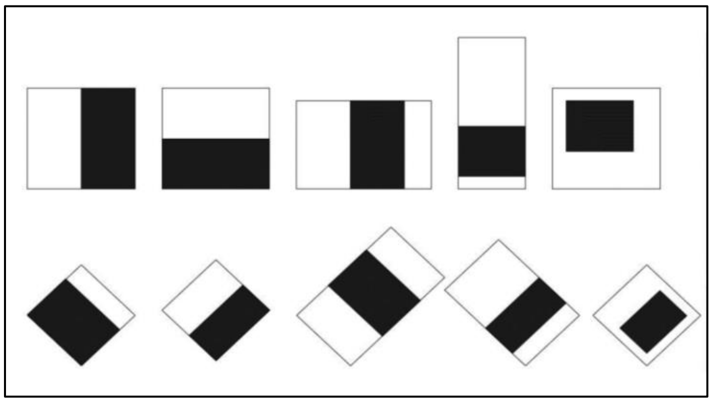

Figure 10 Examples of simple classifiers which may be cascaded for optimal skill and efficiency at run time.

This selection of classifiers is carried out during the algorithm training, where the weak classifier with the lowest error is selected at each step and weight of each classifier is optimized. The algorithm training can take multiple days of computation but it is quick during run time. The training Nortek used incorporates 1855 manually extracted examples of ship shadows and 1671 images without shadows. For the purpose of conceptual understanding, we present a collection of simple classifiers in Figure 10.

There are now three sources entering into the generation of the confidence map (a) the oil-water EMD estimate, (b) land rejection, and (c) vessel shadow detection. Clearly, the beauty of the approach is that it is expandable. It may also incorporate multiple overlaying radar images. It may also have models that use external information such as wind velocity or radar antenna configurations. Figure 11 shows a block diagram of how the confidence map is generated.

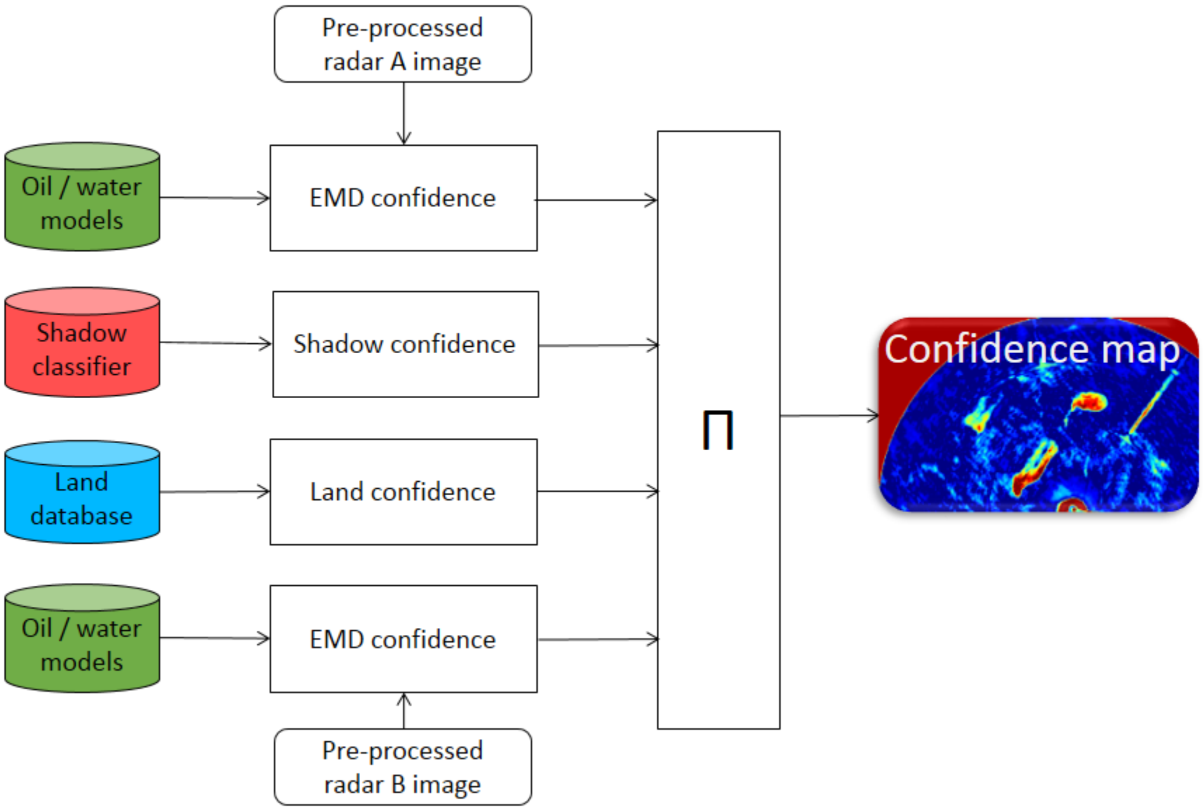

Figure 11 Confidence map generation with different model inputs.

Segmentation

Once the confidence map is completed, the next step in the series for automatic OSD is to evaluate the confidence map for potential oil spills. This step is referred to as "segmentation" and it deals with the association of different pixels with different confident levels, which are neighbors. This is a tiered threshold detector, and effectively associates the higher confidence values with lower confidence values. This does revert back to a form of the threshold detection but now we are operating within the probability (confidence map) domain and not absolute intensity values. The confidence maps have already handled the issue of variability of intensity. The results of the process from raw data to confidence map to segmentation is presented in Figure 12.

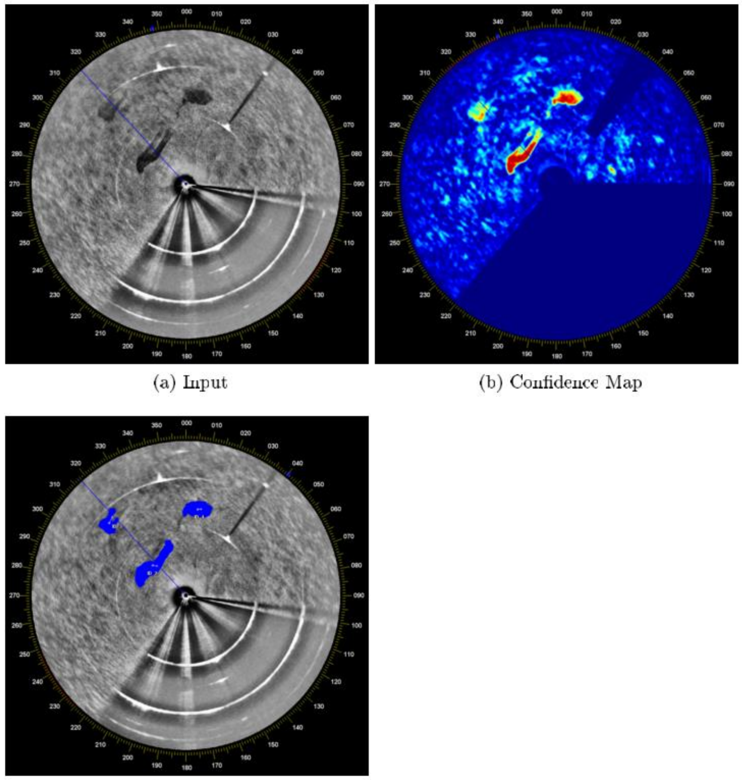

(a) Input (b) Confidence Map

(c) Segmentation

Figure 12 Image examples from the critical OSD steps (a) preprocessed image, (b) Confidence map, (c) Oil Spill Detection result from segmentation.

The remaining steps, which are beyond the scope of this document, are the oil spill tracking and subsequent validation. The tracking is a necessary to associate a detected spill from one sequence image to the next. The time scales are typically minutes or hours long. The ability of the tracker contributes to the overall skill of the OSD. Figure 13 presents an example of the performance evaluation for the OSD. The vertical axis presents detection rate from a scale of 0 to 1, the horizontal axis represents unique oil spills.

The validation step, as indicated in Figure 3, amounts to confirmation that a detection alarm of oil is in fact true or not. This typically requires that an operator access the result. It is becoming more common that this step is aided with an IR camera, which can better access an oil spill once its presence is alarmed by the radar.

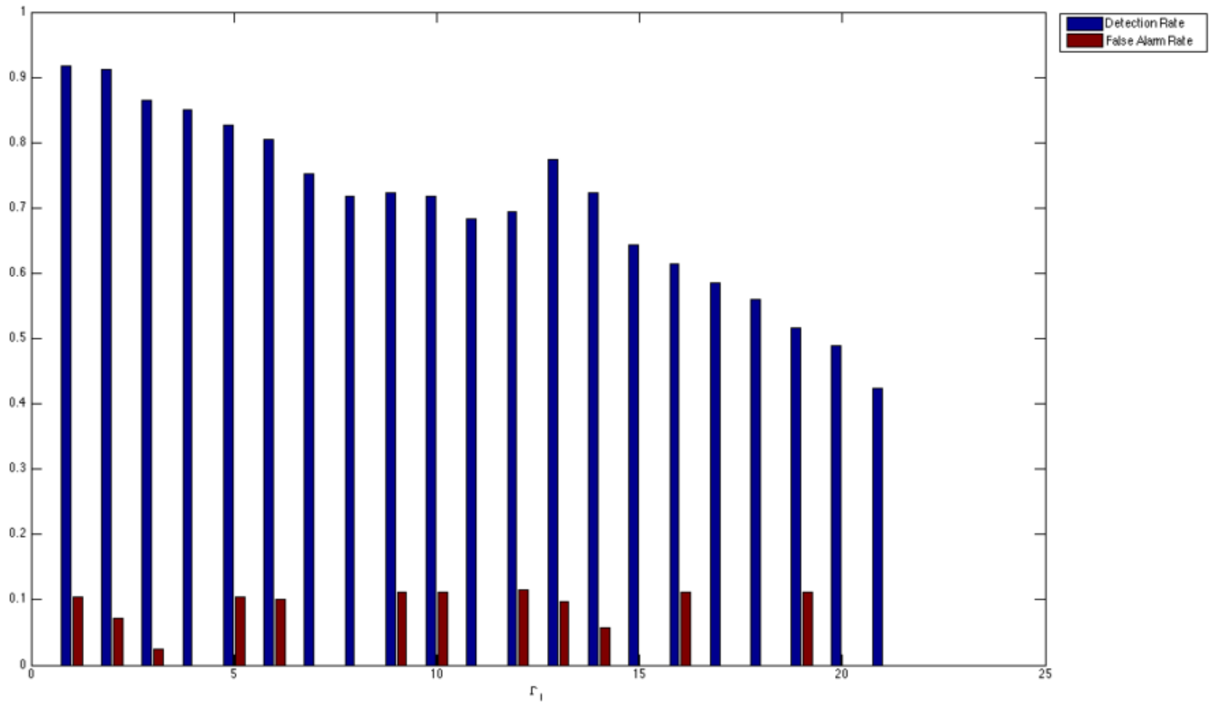

Figure 13 Detection rate (blue) and False Alarm Rate (brown) for 21 unique oil spills.

CONCLUSION:

The methods presented for automatic oil spill detection are those that treat OSD as an image processing problem. The algorithms and models common to this field are implemented at several levels. This includes the Earth Mover Distance histogram comparator for oil characteristics or water characters. The method also includes a shadow detector that employs similar techniques that are used for face recognition in photography.

Apart from being more robust, the model-based approach is also scalable. Multiple data sources may serve as building blocks to establish the likelihood or confidence maps. Complimentary data sources may be winds or sea states, which can modify the comparison models; other sources could also be overlapping radars, which have independent confidence maps. The result is a more enriched and quantifiable indicator for the presence of oil. This leads to greater utility as a remote sensing tool.

The state of the art for OSD radars is one that is moving away from a simple threshold detector and towards model based image analysis. The result benefits stand-alone systems with fewer false detects as well as networked systems which may exchange processed data products for more complete evaluation.

REFRENCES:

Viola, P. and Jones, M., 2001. Rapid Object Detection Using a Boosted Cascade of Simple Features. Computer Vision and Pattern Recognition.

Freund, S., 1999. A Short Introduction to Boosting. Journal of Japanese Society for Artificial Intelligence. 14(5):771-780.

Shirdhonkar, S. and Jacobs, D.W., 2008. Approximate Earth Mover Distance in Linear Time. IEEE Conference on Computer Vision and Pattern Recognition.