

Persistent maritime surveillance using multi-sensor feature association and classification

Sebastiaan P. van den Broek*, Piet B.W. Schwing, Kwan D. Liem, Ric (H).M.A. Schleijsen
TNO, P.O. Box 96864, 2509 JG, The Hague, The Netherlands
*bas.vandenbroek@tno.nl

ABSTRACT

In maritime operational scenarios, such as smuggling, piracy, or terrorist threats, it is not only relevant who or what an observed object is, but also where it is now and in the past in relation to other (geographical) objects. In situation and impact assessment, this information is used to determine whether an object is a threat. Single platform (ship, harbor) or single sensor information will not provide all this information. The work presented in this paper focuses on the sensor and object levels that provide a description of currently observed objects to situation assessment. For use of information of objects at higher information levels, it is necessary to have not only a good description of observed objects at this moment, but also from its past. Therefore, currently observed objects have to be linked to previous occurrences. Kinematic features, as used in tracking, are of limited use, as uncertainties over longer time intervals are so large that no unique associations can be made. Features extracted from different sensors (e.g., ESM, EO/IR) can be used for both association and classification. Features and classifications are used to associate current objects to previous object descriptions, allowing objects to be described better, and provide position history.

In this paper a description of a high level architecture in which such a multi-sensor association is used is described. Results of an assessment of the usability of several features from ESM (from spectrum), EO and IR (shape, contour, keypoints) data for association and classification are shown.

Keywords: persistent surveillance, track recognition, classification, maritime, situational awareness, EO, infrared, ESM

1. INTRODUCTION

Present-day maritime operations take place in a variety of scenarios. In these scenarios a diversity of unwanted activities and threats, such as smugglers, illegal fishery, pirates, or terrorists, may be present. Additionally, the environment is quite complex, containing open ocean, narrow straits, and harbor and coastal environments. Each of these environments has specific background and atmospheric conditions, impacting on the visibility and clutter conditions, in optical, infrared and radar surveillance.

In these operations it is not only relevant who or what an observed object is, but also where it is now and where it was in the past in relation to locations of other objects. In situation and impact assessment, this information is used to determine whether an object is a threat. It is necessary to follow vessels for a long period of time in order to categorize the object and to indicate them as threat. For this assessment it is required that information from different moments is combined. Only at certain times the suspect behavior becomes visible. Furthermore, all the different vessels in the vicinity need to be followed, and the suspect tracks need to be distinguished from normal tracks. Hence persistent surveillance is required.

Single platform (i.e., ship, harbor) or single sensor information will not provide all this information. Fusion of different information sources from different platforms will give a better chance of correctly classifying suspect behavior than single platform-sensor data. The work presented in this paper focuses on the sensor and object levels that provide a description of currently observed objects to situation assessment. For use of information of objects at higher information levels, it is necessary to have not only a good description of observed objects at this moment, but also from its past. Therefore, currently observed objects have to be linked to previous occurrences. Kinematic features, as used in tracking, are of limited use, as uncertainties over longer time intervals are so large that no unique associations can be made. What is required, are object features that are relatively stable over long periods of time. Features extracted from different

sensors (e.g., ESM, EO/IR) can be used for both association and classification. Features and classifications, are used to associate current objects to previous object descriptions from different sensors, allowing objects to be described better, and provide position history. The main algorithmic tasks in our study are hence:

- Fingerprinting: feature extraction from sensor information in order to establish a consistent set of features for each of the detected vessels for object assessment,
- Persistent tracking: correlation of the extracted feature information in order to obtain position history.

In our study approach within the research program Maritime Situational Awareness (MSA) we can build on a large set of feature extraction methodologies developed in the last decade^{1,2,3}. We have reported on feature extraction methodologies based on EO/IR sensors making use of global descriptions as well as local descriptors. One local approach is based on keypoints, discriminating between jet-ski, RHIB (Rigid Hull Inflatable Boat) and Watertaxi. ESM data has been analyzed, and demonstrated a successful persistent tracking over a three month period⁴. These ESM and EO/IR datasets contain complementary data that can be applied successfully in feature fusion techniques.

In this paper we present our approach to persistent surveillance in the research program MSA⁴, with an emphasis on the work done on multi-sensor track recognition. In Section 2 we present the high-level architecture in which such a multi-sensor association is used. Section 3 contains an assessment of the usability of several features from ESM and EO/IR data for association and classification. Section 4 contains discussions and conclusions.

2. ARCHITECTURE

The persistent surveillance research is part of a larger research program on Maritime Situational Awareness that further includes research on higher level information fusion and command and control. That information fusion uses information obtained for the observed ships from different sources to determine intent. This is used in the command and control for long and short term platform planning. Platform movements then cause new sensor readings and these create updates of object information, closing the loop. Figure 1 shows a high level view, with information level increasing from sensor level at the left, via object assessment to situation assessment and platform management. Besides research performed using real data for the lower information levels, this information fusion chain is also being implemented in a simulation framework.



Figure 1. High level MSA architecture view, showing information processing at different information levels.

The main functionality developed for object assessment consists of track recognition and object classification (see figure 2). Here, track recognition is defined as recognizing a currently tracked ship to be the same as a previously observed track. Object classification is the description of a ship as belonging to several categories. These two functionalities will be combined into a single processing module, which is used for:

- Processing real data: In this case, the processing is applied on objects with features that are derived from real data, as it would also be in operational use.
- Simulation: In this case, as part of the simulated chain, no sensor data is available, but measurements of objects with features are simulated, for demonstration and algorithm assessment.

In addition to the track recognition and classification algorithms, for the real data chain, algorithms are being developed to extract features from sensor data. For the simulated chain, models need to be defined to allow simulation of the features, based on simulated objects, sensors and environment. The simulation chain allows demonstration of the use of track recognition and classification in more complex scenarios than recorded data allows, and is further aimed at developing the higher level algorithms such as situation and intent assessment. Object detection and kinematical tracking are not part of the current research. Algorithms for these are available from earlier work⁵ in near real-time implementations.




Figure 2. More detailed view of the object level processing. Extracted features are used both for track recognition and classification.

With small update intervals, kinematical tracking is sufficient to predict the location of previously observed ships accurately enough to limit the association problem (from observed ships to existing tracks) enough to be able to solve it. In this prediction from last track estimation to current observation time, often a simple model for movement of the objects is assumed, such as a constant velocity or a constant acceleration model. Although the ship may not move according to the model, deviations from the assumed model can be accounted for in a modeling error. For longer intervals, this deviation causes the predicted area to be too large to be useful, with respect to the number of ships present. As a result, observed ships could be associated to many tracks, causing problems finding a correct association. For persistent tracking of ships over hours, or days, the association problem should be reduced. This can be done on two sides: reducing the area in which tracks are predicted, and expanding the number of features on which association is done beyond position. For reduction of the area, as a minimum, a maximum speed could be used to limit the distance a ship could have traveled, assuming a shortest (straight) route. Further improvements could take into account actual geography, to compute an actual shortest route, and other boundary conditions as shipping lanes etc. Taking more than position into account in association, allows for reducing the number of possible associations, by calculating the likelihood that a newly observed ship is the same as a previously observed one. This process is performed by removing associations with high similarity for which a match is certain, and by excluding associations when similarity is very low. In between, similarity can be used to make the most optimal choices. The main topic of the current research is in finding features that are descriptive and/or discriminative enough, with long term stability, to be used as similarity measure, and methods to use this information in association and track recognition.

All features used for track recognition can also be used in object classification, where objects are assigned to different categories. These can be different types of ships (e.g., fishing trailer or tanker), but also other divisions, such as based on size (small, large). Such categorizations are used at higher level processing to determine intent.

3. ASSESSMENT OF FEATURES FOR PERSISTENT SURVEILLANCE

In this section we present an assessment of useful features for fingerprinting in recognizing ships. The focus of current research is in the use of electro-optical (visible light and infrared) imaging, and specific emitter identification using ESM (Electronic Support Measures) data, which will be presented in the next sections. In this way the approach is AIS (Automatic Identification System) independent, and an addition in case AIS is not available or cannot be trusted.

3.1 Electro-optics

In earlier work, several features extracted from electro-optical data (with focus on infrared) were examined, for classification and clutter reduction^{2,3,6}. This evaluation of features, or optical fingerprints, and classification methods also

gives an indication of descriptiveness (towards class) and separation between more general types. These are also an indication for use in track recognition. Although the compilation of an operational picture requires persistent surveillance of all types of vessels, detailed evaluations were performed mostly on small ships, because of data availability. However, these experiments do allow a more qualitative assessment based on properties observed for the different features. Below, evaluations of several features are discussed: moments, contours, parameterized descriptions, other global descriptions, and local keypoint features.

Moments

Moments give a global description of how information (detected pixels, intensity values) are distributed over the observed area of an object. Higher order moments for example give an indication of horizontal or vertical symmetry, or whether the shape is concentrated around the center or near the outside. We used normalized moments⁷ for discriminating between classes of small ships before⁶. It was shown on infrared recordings of targets such as a small cabin boat, RHIB and jet-ski, that a good distinction is possible, even when the observed detection of the objects is only a few pixels high, and with no detailed information on the objects visible. Figure 3 shows examples of moments for the three ships. The determined moments depend on the aspect angle at which the object is observed. Clearly, as frontal views are more symmetric and compact, higher order moments will be smaller and less discriminative for most ship types, as compared with side views, where differences will be larger. By choosing a good combination of moment orders and using a Mixture of Gaussians classifier², a classification probability of the correct type of over 80 % was obtained for side views (for which the classifier was trained), but even for frontal views some distinction could be made. By combining classifications over time, the results improve.




Figure 3. Scale invariant moments for RHIB (-), water-taxi (--) and jet-ski (-), at distances of 500 and 1000 m, for side views. The moment orders (0,2), (0,3) and (3,0) are shown.

Qualitative observation of moments for silhouettes of larger ships such as container vessels and sailing ships showed that at least a global distinction between such types of objects is possible.

In track recognition, comparing the moments for current tracked objects with earlier observed ship descriptions, may not allow for a unique identification, as moments are not that descriptive. However, it will allow for a good indication of whether they could be the same, and more importantly, exclude many possible associations.

The information that determine the moments depends on aspect angle, but also distance has an effect due to resolution of the camera (less pixels on target at larger distance) and due to atmospheric effects (transmission and blur over longer range). Knowing the distance to the target (for example obtained from radar) and the aspect angle (e.g. estimated from radar track course) will help to determine what to expect from moments, when comparing to other features.

Contours

Height contours of segmented ship detections can contain detailed information about a ship. Examples of segmentations are given in figure 4. To obtain a useful contour (the height of the segmentation from side to side) a good segmentation

of the detection of the ship is needed. As the figure shows, contours can be quite descriptive, in indicating masts, chimneys, cabins, relative size of hull, etc. However, aspect angle affects the view. This effect is not as big as it may seem, as it was shown that by normalizing the width of the contour, high correlations between contours are obtained, for most aspect angles, except for about 20 degrees around the frontal view⁶. Observation distance also influences the usefulness of using contours, as with lower resolution details are lost, and correlations reduce. Finally, atmospheric effects, causing loss of contrast and distortion, make obtaining a good segmentation difficult. Especially smaller details will be lost, for smaller ships or larger distances, hence the advantage of using contours over using for example moments is smaller, as the contour information will be less descriptive. In these cases, guidance input from an operator may be needed.



Figure 4. Silhouettes of a tanker at different aspect angles.

When detail is high (at close enough range and under non-adverse atmospheric conditions), contours allow for a good association between observed ships and earlier recorded contours of tracked vessels. Even contours with somewhat lower quality are discriminative and will greatly reduce the number of possible associations.

Parameterized description

Correlations between contours are strongly depending on small changes in estimated location of vertical changes, such as edges of cabins and masts, as a result of limited resolution or atmospheric distortions. By parameterizing the description of distinctive features, such as chimneys, cabin, hull, etc., matching can be improved, at the cost of losing some distinctiveness in the description. A description can be obtained from a contour by grouping parts with similar height, e.g., using a (height value) histogram², as indicated by the example in figure 5. The resulting description will not be enough to identify a specific ship (as may be possible with a good contour), but it is more discriminative than moments, as it eliminates possible associations based on larger differences. As with contours, at (near) frontal views, most distinctive details will not be visible.

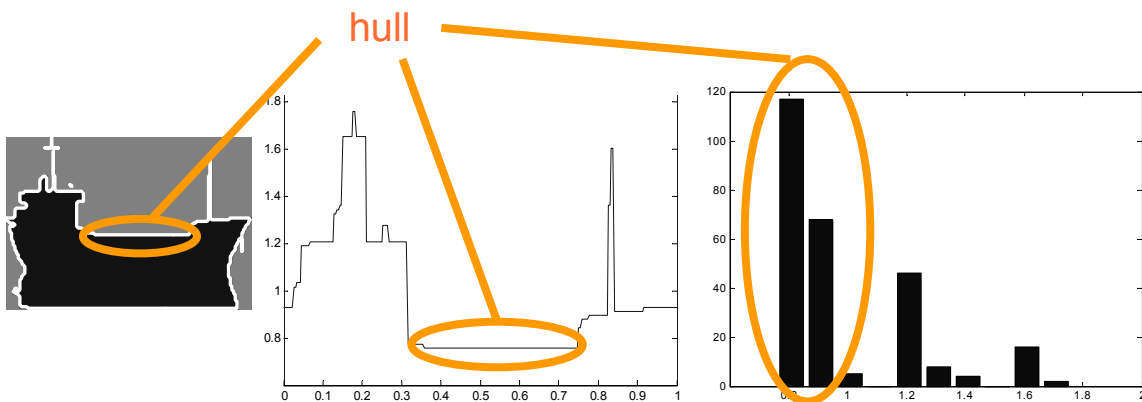


Figure 5. The middle graph shows the height profile of the silhouette of the ship shown on the left. On the right the histogram of height values of this profile is shown. Groups in the histogram indicate hull, superstructure and masts.

Other global properties

When distance and aspect angle are known (e.g., from radar measurements), width and height of the ship may be estimated and used as a first simple distinctive feature. At larger distances however, resolution limitations and segmentation problems may make these values less accurate and hence less discriminative. From width and height an aspect ratio can be determined, even when distance is unknown. Another feature is the compactness, defined as the ratio between the area and circumference of a segmented detection². It gives an indication of how intricate the shape of a ship is. As it discards most of the information contained in a contour about where specific features are, but does describe in some extent the existence of chimneys, cabins, masts, etc., it is a simple feature, in quality in between contours and simple measurements of size. Although it has an aspect angle dependence, it is less sensitive to resolution and atmosphere than contours. It will not be as descriptive, and therefore not allow as unique an association as contours provide, but it is discriminative enough to help in reducing the possible associations. Track information gives an indication of aspect angle (estimated from the course of the vessel), allowing for some correction of the aspect angle dependence.

Local keypoint descriptors

Keypoints such as SIFT⁸ are commonly used in image recognition, and for example image stitching. In these applications, the images are generally high resolution color images with much detail. The keypoints describe a local image area based on for example intensity or intensity gradients, at locations where such information content is high (e.g., corners). Application of such methods to imagery of ships in maritime surveillance appear to be less useful, as detail in images is generally lower. This low contrast can be due to atmospheric effects when observing at larger distances, or less intensity variation in general, for example in infrared, where apparent temperature variations are typically less detailed as visible light features. However, our research^{2,3} showed that the use of keypoints in discriminating between small ships in imagery at infrared wavelengths is possible. Best results were obtained using a Harris corner detector⁹ to determine descriptive locations, in combination with either a SIFT description, or a local region description. Figure 6 shows examples of descriptive keypoint locations in two infrared images of small vessels. An advantage of the use of keypoints is that aspect angle dependence is reduced, as each of the keypoints itself is descriptive and if some are not visible, a match can still be made. Also, even though in infrared much detail is found on the edges of the ship, sensitivity to segmentation errors is smaller. Using the relative locations of keypoints allows for making an overall description of the ship. At lower contrasts and resolutions, this allows for distinction between types of ships, even with less intensity detail such as in infrared. At good resolutions the descriptions can be used as a similarity measure for recognition.




Figure 6. Infrared images of Jet-ski (left) and water taxi (right) with detected keypoints.

3.2 Specific Emitter Identification using ESM sensors

The majority of navigation radars have pulsed emission characteristics which are transmitted as pulse trains. Current Radar ESM sensors detect and measure the characteristics of individual pulses in these pulse trains. Typical pulse

features are time of arrival, carrier frequency, pulse width and direction of arrival. These parameters are the main input data for the separation and reconstruction process of the original pulse trains in the ESM sensor. The derived inter-pulse parameters of the reconstructed pulse train data, like mean RF (frequency), mean pulse repetition interval and pulse to pulse parameter variations, are then compared with the content of a database of known pulse train data records of radar emitters of interest, in order to determine the identity (ID) of these radars. Typically the ESM ID result is reported at a radar equipment type level, e.g. Decca navigation radar type X.

Incorporation of digital receivers with high signal sampling rates in new generation ESM sensors provides these sensors with an additional capability to derive and store the details of the signal modulation characteristics within each individual radar pulse, the so called intra-pulse modulation characteristics. The presence of fixed intended and unintended modulation patterns in the pulses of particular radar transmitters can then be used for Specific Emitter Identification (SEI)^{10,11,12,13,14}. SEI is exploiting these intra-pulse patterns as electronic fingerprints to recognize and identify particular radar transmitters among radar transmitters of the same equipment type, e.g. the Decca navigation radar type X on ferryboat A. The SEI result can be used for different purposes, e.g., to enhance the ESM sensor contribution to situational awareness and for long term monitoring of a particular ship in crowded coastal areas⁴. The currently chosen SEI method in our experiments is based on PCA (Principal Components Analysis).

SEI Method description

Example of feasibility of discrimination of navigation radar transmitters through their intra-pulse modulation patterns is shown in the upper part of figure 7, for three different parameters: intra-pulse amplitude variation, intra-pulse RF (frequency) variation and intra-pulse non-linear phase variation.

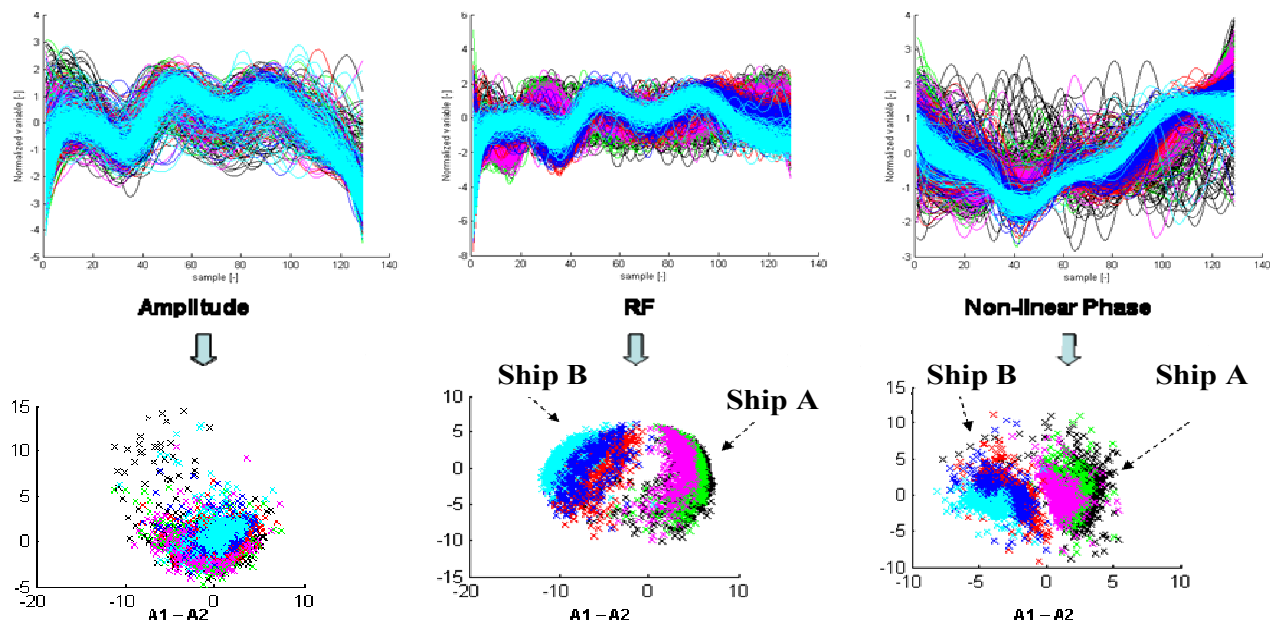


Figure 7. Result of PCA projection of the intra-pulse patterns of Ship A radar pulses (± 2000) on Ship B radar pulses (± 1200), collected over a period of 3 months, into the 2D feature space spanned by the first two principal components axes. Upper part: stacked raw intra-pulse patterns. Lower part: corresponding projections in the PCA space. Color coding is applied to indicate the source and the collection time (month) of the data.

The shown pulse data patterns (top part of figure 7) belong to two navigation radars of the same type which are installed on two ships: ship A and ship B. The data are collected over a period of three months using PESMO, TNO's wideband digital receiver with a sampling rate of 1 G sample/s and 8 bit AD resolution. Over the whole period of three months, two different groups of stable intra-pulse modulation patterns are distinguishable for the RF and the Non-linear phase intra-pulse parameter, each belonging to the radar of one of the two ships. Each intra-pulse pattern of a single radar pulse

here can be considered as a data vector with 130 elements. The PCA transformation is then applied on these data sets in order to reduce the amount of data needed to describe the intra-pulse patterns, through automatic derivation of the first few principal components that account for most of the variation in the original data, with little loss of information. Similarities and dissimilarities between the intra-pulse features of radar pulses can then be emphasized using smaller feature data sets.

The lower part of figure 7 shows the result of projection of the complex intra-pulse patterns of the two radars into a PCA feature space spanned by the first two principal component axes. In this PCA feature space the similarities and dissimilarities between the complex intra-pulse patterns of the individual pulses are translated into distances between the projections of these patterns. Despite the small dimensions of the feature space, the two clusters of points that represent the unique RF and non-linear phase of the pulses of the two transmitters can already be clearly distinguished.

Functionally, the SEI implementation consists of an initial SEI training phase and a subsequent automatic SEI matching phase. In the training phase a representative set of measured intra-pulse patterns of the emitters of interest is selected, to function as the PCA training set. Based on the intra-pulse data patterns of the radar transmitters in the PCA training set, a PCA transformation matrix is calculated. During the automatic SEI matching phase, the measured intra-pulse patterns of the unknown pulses, together with the intra-pulse patterns of the PCA training set, are projected into the PCA feature space using this PCA transformation matrix. The ID of an unknown pulse is then automatically determined using the k -nearest neighbour classifier, depending on the distance of its projection point in the PCA feature space to the ID labelled projection points of the pulse patterns in the PCA training set. The most occurring training set ID label within the k -nearest training pulse projections is then assigned to the unknown pulse, where k is a user defined value.

Table 1. Example of PCA SEI performance using Non-linear Phase intra-pulse patterns. Overall Specific ID performance of 82 % is the average of the diagonal elements. Total test set: 11000 pulses of 7 navigation radar transmitters on 6 different ships (A to F) collected over 3 months period. The radars are products of three different manufactures (X to Z). PCA feature space has 5 dimensions. The chosen k in the nearest neighbor classifier is 5.

		Input data source						
		Ship A Rdr X	Ship B Rdr X	Ship C Rdr Y	Ship D Rdr Y	Ship E Rdr Z	Ship F Rdr Y1	Ship F Rdr Y2
Reported ID Result	Ship A	93.1%	0.5%	0.0%	0.3%	0.5%	0.2%	0.3%
	Ship B	2.0%	77.3%	0.0%	3.9%	0.9%	0.2%	0.3%
	Ship C	0.8%	3.3%	97.9%	3.0%	0.0%	38.5%	0.9%
	Ship D	1.4%	0.1%	1.2%	67.4%	0.5%	3.5%	0.3%
	Ship E	0.6%	0.5%	0.9%	5.5%	98.0%	0.5%	0.0%
	Ship F_Y1	0.0%	13.0%	0.0%	0.2%	0.0%	47.6%	0.3%
	Ship F_Y2	2.0%	5.2%	0.0%	19.9%	0.0%	9.6%	97.9%

Results on SEI

Table 1 gives the observed performance of the PCA SEI method for 7 different navigation radars, using the Non-linear phase intra-pulse of a set of 11000 pulses that are collected during 3 months. The resulting overall SEI performance is 82 % correct results. For amplitude and RF intra-pulse patterns, the PCA SEI method gives lower correct results, respectively 75 % and 76 %. Note that in case of ship F there is uncertainty about the number of the installed radars on

this ship. The PCA cluster data shows clearly two different clusters (Y1 and Y2) while a check in other sources only indicates the presence of one single radar on ship F. Where ships are not recognized correctly, combination with data from other sensors (such as EO) will help.

Application of other methods to reduce the amount of data in the intra-pulse radar signatures, that deliver less abstract information (currently “points in PCA space”) than the statistical PCA method, may be needed to detect slightly different fingerprints of unknown radars that are not yet included in the PCA training data set. An example of such methods is application of curve fitting methods to deliver a parametric description of the intra-pulse fingerprint with less data, e.g. polynomial curve fitting methods.

The SEI information allows for specific matches of newly observed ship radars, with stored information about earlier observations. This match enables a coupling to (more descriptive) information of that ship available earlier, such as classification from EO information, or even the name of the ship. In practise, the distance at which SEI information can be collected is only limited by the radio horizon. The shown SEI results are collected at a distance of 25 km from the ships. SEI information is less dependent on the aspect angle at which the ship is observed because navigation radars have rotating antennas. Therefore in practise the radar antenna main beam will always periodically illuminate the ESM sensor location in the azimuth plane.

4. DISCUSSION AND CONCLUSION

For higher level information fusion for situation assessment a history of observed ships is needed. At larger observation intervals kinematical tracking will no longer be sufficient to link observations in time. Gating and association solely based on position should be extended with looking at similarity based on other features. In the previous section, several such features were described, as obtained from EO/IR imagery, and from ESM. General observations on the usability of the features from this assessment are summarized in table 2. Indicated is whether features are better (+) or less (-) suited for use in recognizing a specific ship (uniqueness) or in distinguishing between types of ships (helping association and object classification). Dependencies on distance, aspect angle and atmosphere are for both quality and stability. For example, for contours, both quality of information, and the description provided by the contour change with distance.

Table 2. Summary of assessment of features for uniqueness (recognizing specific ship), distinctiveness (distinguishing types), and dependence on distance, aspect angle and atmosphere.

Feature	Uniqueness	Distinctiveness	R-dependence	Aspect dependence	Atmospheric dependence
Moments	-	+	0	-	0
Contours	+	+	-	-	-
Parametric descriptions	0	+	-	-	0
Global properties	-	0	0	0	0
Keypoint descriptions	+	+	0	0	0
ESM SEI	++	-	+	++	+

Each has its advantages and disadvantages due to dependence on observation aspect angle, distance (resolution) and atmospheric influences. In this way, some are quite descriptive and discriminative and can be used to recognize ships directly. Others are merely distinctive between categories of ships, and mostly increase similarity relative to other possible matches by excluding possibilities. It therefore makes sense to combine features and sensors. To do this, some weighing is needed between similarities based on different features, determined for the current accuracy (due to for example resolution limitations) and certainty (i.e., expected correctness at current distance and atmospheric conditions).

When several observations are done during an interval that an object is tracked (for example by cueing a camera on a radar track, see figure 8), an estimation of the variation is available as an indication of the accuracy of the derived similarity.

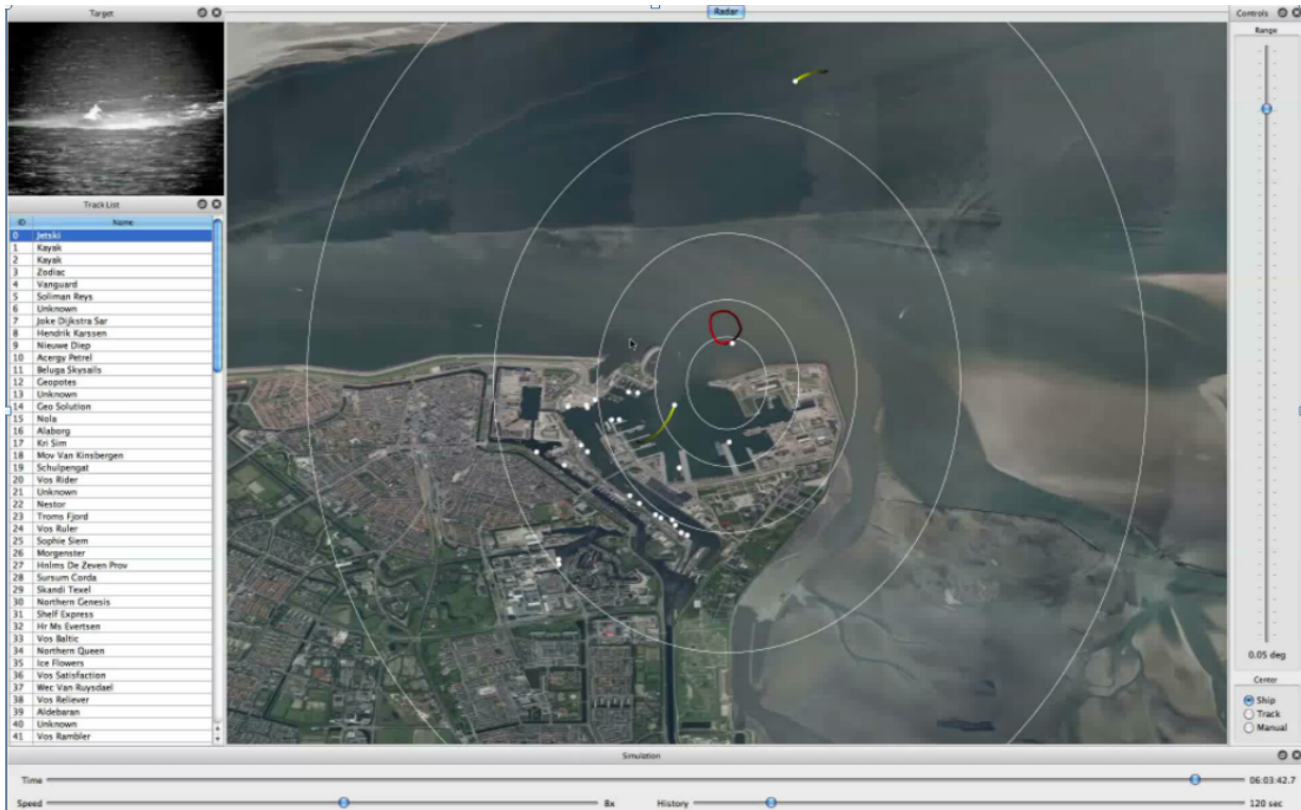


Figure 8. Example of a possible view of tracks with cued camera recordings for track recognition and classification.

At closer ranges, (EO/IR image) contours may allow for quite specific recognition. At larger distances or adverse atmospheric conditions, moments and more simple features such as compactness are a better choice. Parameterized descriptions (such as number of masts) are a good compromise, being descriptive and distinctive. Local keypoint descriptions are robust and usable at close range for recognition, and at larger distances for discriminating between types of ships. Specific emitter identification from intra-pulse ESM signatures is less descriptive of the type of ship, but is quite descriptive of the specific radar, and is a good feature for recognition, over long time intervals (of several months). A big advantage of ESM features is that they are not depending on observation aspect angle, and can be measured at much larger distances than EO and IR features. Clearly, the combination of ESM and EO/IR will combine the advantages.

In a processing chain for situational awareness, as is the main topic of the MSA research program, information at different levels is used to improve the track continuity. For example, sensor and environmental information provide indications of the usefulness of different object features in the categorization process. For example, infrared target imagery for the current environmental conditions can be modeled¹⁵ and provide an estimation of image quality such as contrast and clutter.

ACKNOWLEDGEMENTS

This research was supported by the Netherlands MOD under the program V1114, Maritime Situational Awareness. The authors wish to thank F. Bolderheij (NLDA) and the MSA program team for the valuable discussions and input.

REFERENCES

- [1] Schwering, P.B.W. , Lensen, H.A., Broek, S.P. van den, Hollander, R.J.M. den, Mark, W. van der, Bouma, H., Kemp, R.A.W., "Application of heterogeneous multiple camera system with panoramic capabilities in a Harbor Environment", Proc. SPIE 7481, 74810C (2009)
- [2] Broek, S. P. van den, Bouma, H., Degache, M. A. C., Burghouts, G., "Discrimination of classes of ships for aided recognition in a coastal environment", Proc. SPIE 7335, 73350W (2009)
- [3] Mouthaan, M.M., Broek, S.P. van den, Hendriks, E.A., Schwering, P.B.W , "Region descriptors for automatic classification of small sea targets in infrared video", Opt. Eng. 50, 037002, doi:10.1117/1.3549907, (2011)
- [4] Smith, A.J.E., Anitori, L., Bergmans, J., Colin, M., Iersel, M. van, Liem, K.D., Schwering, P.B.W., Sweeden, R. van, Vullings, H.J.L.M., "Overview of maritime situational awareness research at the Netherlands Organization for Applied Scientific Research TNO", in 'NATO Workshop MSA 2009', 15-17 September 2009, La Spezia, Italy, (2009)
- [5] Bouma, H., Lange, D.J.J. de, Broek, S.P. van den, Kemp, R.A.W. and Schwering, P.B.W., "Automatic detection of small surface targets with electro-optical sensors in a harbor environment," Proc. SPIE 7114, 711402 (2008)
- [6] Broek, S.P. van den, Bouma, H. and Degache, M.A.C., "Discriminating small extended targets at sea from clutter and other classes of boats in infrared and visual light imagery," Proc. SPIE 6969, 69690B (2008)
- [7] Sonka, M., Hlavac, V. and Boyle, R., "Image Processing, Analysis, and Machine Vision," Brooks/Cole Publishing Company, (1999)
- [8] Lowe, D.G., "Distinctive image features from scale-invariant keypoints," International Journal of Computer Vision 60(2), 91-110 (2004)
- [9] Harris, C. and Stephans, M., "A combined corner and edge detector," Proceedings of the 4th Alvey Vision Conference Manchester, 189-192 (1988)
- [10] Talbot, K.I., Duley, P.R. and Hyatt, M.H. , "Specific emitter identification and verification", *Technology Review Journal*, (2003).
- [11] Langley, L.E. , "Specific emitter identification (SEI) and classical parameter fusion technology", *WESCON 1993, Conference Record*; (1993).
- [12] Eilevstjonn, J. , Odegard, J.E. and Malnes, E. , "Radar pulse classification using compressed intrapulse feature vectors", *9-th DSP Workshop 2000*, (2000).
- [13] Luo, Z.-Q., Liu, J., Lee, J., Wong, K.M. and Gao; S. , "Radar emitter classification using intra-pulse data.", *International Journal on Electronics and Communications*, Volume 53(Number 6), pp. 324-332, (1999).
- [14] Kawalec, A., and Owczarek, R., "The method for radar signal recognition", *Journal of Physique IV France*, Vol. 137, (2006).
- [15] Kunz, G.J., Degache, M.A.C., Moerman, M.M., Eijk, A.M.J. van, Neele, F.P., Doss-Hammel S.M. and Tsintikidis, D., "Status and Developments in EOSTAR, a model to predict IR sensor performance in the marine environment," Proc. SPIE 5572, 101-111 (2004)