Discriminating small extended targets at sea from clutter and other classes of boats in infrared and visual light imagery

Sebastiaan P. van den Broek*, Henri Bouma, Marianne A.C. Degache
TNO Defence, Security And Safety, P.O. Box 96864, 2509 JG, The Hague, The Netherlands
*bas.vandenbroek@tno.nl; phone +31 70 3740442; fax +31 70 3740654

ABSTRACT

Operating in a coastal environment, with a multitude of boats of different sizes, detection of small extended targets is only one problem. A further difficulty is in discriminating detections of possible threats from alarms due to sea and coastal clutter, and from boats that are neutral for a specific operational task. Adding target features to detections allows filtering out clutter before tracking. Features can also be used to add labels resulting from a classification step. Both will help tracking by facilitating association. Labeling and information from features can be an aid to an operator, or can reduce the number of false alarms for more automatic systems.

In this paper we present work on clutter reduction and classification of small extended targets from infrared and visual light imagery. Several methods for discriminating between classes of objects were examined, with an emphasis on less complex techniques, such as rules and decision trees. Similar techniques can be used to discriminate between targets and clutter, and between different classes of boats. Different features are examined that possibly allow discrimination between several classes. Data recordings are used, in infrared and visual light, with a range of targets including rhibs, cabin boats and jet-skis.

Keywords: Classification, clutter reduction, target recognition, small targets, boats, jet-ski

1. INTRODUCTION

Tasks of naval forces are more frequently taking place in coastal environments. Such environments introduce specific challenges in building an operational picture for situational awareness. For example, clutter may be caused by breaking waves and coastal backgrounds, and may be more irregular than at open sea. Also, ships that are encountered not only include larger vessels, such as merchant boats, fishing trailers and frigates, but especially smaller boats, such as cabin boats and jet-skis. Finding those ships that are of interest for an operational task (such as intercepting smugglers), or those ships that could pose a threat, is a challenging task. Performing those tasks is often done by a human operator, based on radar information and using electro-optical sensors for final identification. The operator may be helped by a system that adds information from electro-optical sensors automatically, assisting him in selecting targets for further inspection, and helping in classifying ships as targets of interest. The process of finding targets of interest can be defined as separate steps, from sensor data acquisition and detection to classifying a target as a threat or smuggler. Some of these steps may be done by an operator, some automatically, adding information for the operator to use, by labeling detected objects.

At TNO a research programme “Electro-optical sensor systems” for the Dutch ministry of Defence is examining several aspects of the use of electro-optical systems, from (near)-future system concepts to the image processing chain [12]. Work on image enhancement is described in [8]. The project “Classification and clutter reduction” is part of this programme and examines ways of adding information in the process to assist an operator in classifying ships, with an emphasis on small targets in a coastal environment. This includes what information can be extracted, and how this can be used to discriminate between different classes, at different levels. Examples of such levels are discriminating between objects or clutter, big or small ships, less interpretable classes (such as flat or tall, (a)symmetric, intricate shape or not), or type of ships (such as frigate, merchant vessel, fishing trailer, cabin boat, jet-ski). Whether such a discrimination is possible, depends on many things. Which features can be determined that are of use for discriminating between classes? Which methods for automatic discrimination can be applied? How to handle the use of features in a database, or the lack thereof? In this paper we present ongoing work in this project. A general description of a classification process is given,

Signal and Data Processing of Small Targets 2008, edited by Oliver E. Drummond
Proc. of SPIE Vol. 6969, 69690B, (2008) · 0277-786X/08/$18 · doi: 10.1117/12.777542

consisting of several steps where features are determined and used to label detections. Examples are given how features allow discrimination between targets and clutter, and of discrimination between small extended targets from infrared and visual light imagery. The influence of the aspect angle (defining how the object is seen by the sensor) is discussed, and it is shown how knowledge of the aspect angle makes features more useful in the classification process.

The process is discussed in the next section, In Section 3, results for clutter reduction are shown. In Section 4, the dependency on aspect angle of height profiles is discussed. This dependency is also apparent when looking at small targets, as shown in Section 5. In this section, several features are computed and several classifiers have been applied to discriminate between small targets. Finally, a conclusion is given in Section 6.

2. DESCRIPTION OF PROCESS AND DATA

Figure 1 shows a representation of the process of adding information to detected objects. Several steps add information. The segmentation step is the link from detection to the rest of the process, as the shape of the detection is of more importance than it is in the detection step itself, where mostly only contrast is used. The work described in this paper relates to the steps in the process between detection and showing object information to an operator. The column on the left of the flow describes a process that runs continuously in a loop. It describes an automatic process for detection in electro-optical sensor data, which could be used instead of radar, or (more often) in addition to radar tracks (in a multi-sensor process). In automatic detection there is a trade-off between detection probability and false alarm rate. Especially when interested in small boats in a coastal environment, false alarms cannot be avoided without compromising detection results. When a process is used to filter away false alarms, a higher detection probability for targets of interest can be achieved, with fewer false alarms. Clutter reduction is discrimination between targets of interest and other alarms. Tracking can also be used for filtering false alarms, by using information in time, such as track lifetime or temporal behavior of features. Other methods of labeling detected objects as belonging to one of different classes are used to add information to the tracks.

The right side of the figure shows less frequent labeling of objects with information, as it is not done for each image. This includes more complicated processing, and is also the part of the process that may involve an operator. For example it can provide information that helps an operator in making a decision on the nature of the target. When linked to tracks, it can make use of time dependent features. It improves long time tracking, by making association over time more robust. Other sensors may be involved in the process, for example radar tracks could trigger the process to obtain information from an image for that track. Data from other sensors can also help by adding more accurate distance and aspect angle information to the features computed from the electro-optical data.

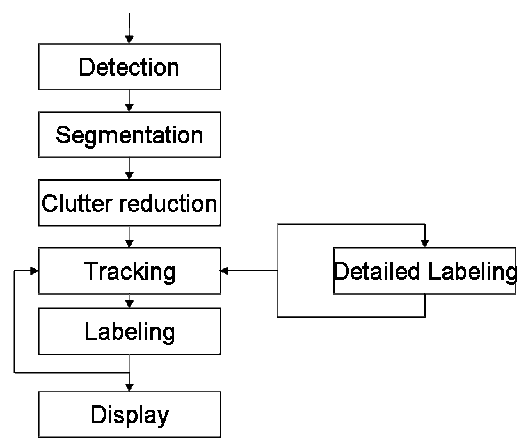

Fig. 1. Process of adding information to detected objects. Input are images, all other lines indicate object information (including segmentation, features, track information, labels from labeling)

Several of the functions described in Figure 1 include the computation of features, and using features to label an object as belonging to one class or another, or assign probabilities of the object belonging to classes. This selection can be made using rules determined by hand, or automatically learned. The latter should be done with care, checking that the system did train on the relevant information. In literature, for classification of ships (or similarly aircrafts) often more complex features describing the extracted shape are used. Examples are moments [1, 4, 10, 11] and Fourier descriptors [11, 14]. Methods for training include Principal Component Analysis [6] and neural networks [1, 10]. We examined the use of moments, with respect to aspect angle dependence and for use in discriminating targets that have segmentations that consist of much less pixels than generally used.

In order to discriminate between classes based on features, image data with relevant information is needed to find how features define the classes. Often the number of objects in available data sets is limited. In other words, many targets of interest are not covered by the examined data. Also, in many cases not enough information about the data is known to fully exploit it. Ideally, everything affecting the features should be known, e.g., type of ship, location (distance), aspect angle, atmospheric conditions, and waveband. As this is not always the case, this often poses a problem when researching possible features and methods, but will also be a problem in practical applications. Automatic training may lead to unexpected results. The fact that not all possible targets are in a training set, also limits the use of an automatic classifier. There should be enough intelligence in the process to define a class that contains the targets not covered by the defined classes, so they can be labeled as unknown.

Table 1. Data used for algorithm development and testing

No.	**Name**	**Description**	**Target information**
1	SPITS	Visible light, single images, many ships	No target information
2	SURFER	3-5µ and visible light sequences. Several small targets	Targets, distance, sea state, meteo.
3	Rotterdam	3-5µ, 8-12µ, visible light. Sequences. Harbor, small known targets and targets of opportunity.	3 known targets (with some distance information), and unknown targets.
4	MCG/8 Stavanger	8-12µ. Frigate.	One known target with accurate distance and aspect angle information.
5	Simulated frigate data	Simulated silhouettes of a frigate of same type as 4	Known target (not entirely correct model). Known distance and aspect angle and meteo.
6	SPITS simulation	Simulated IR images, three targets	Known targets with distance and aspect angle (see also [2])

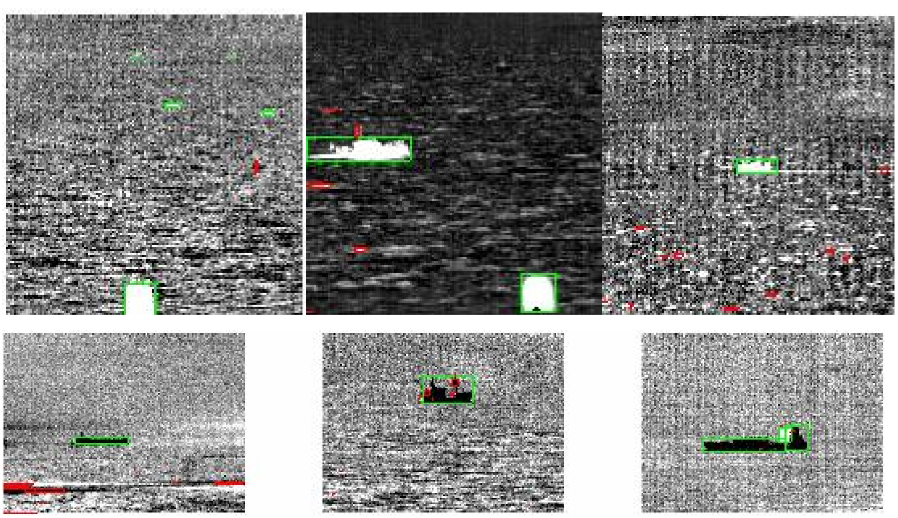

Fig. 2. Example of some of the data sets. Shown is contrast scaled to local background variations. Rectangles indicate automatic detections of targets and background. Top images are from IR recordings, bottom from visual light.

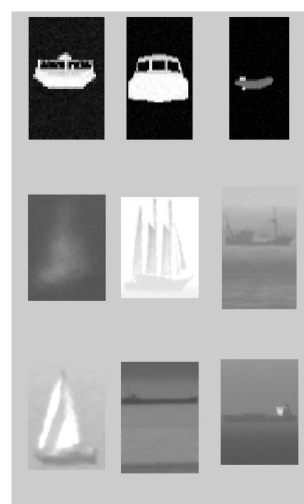

Fig. 3. Example of targets from data sets 1 and 6. The top three are simulated using EOSTAR [9].

Limited data sets can be used to examine parts of a process for discriminating between classes, to get an idea of usability of features and methods, and to examine dependencies on parameters such as distance and aspect angle. Artificial data, simulated from models of ships, may be used to augment available data, or fill in gaps in databases. We simulated data using TNO software EOSTAR [9]. Data used in our research is described in Table 1. Some examples of the different data is shown in Figures 2 and 3, including simulated data of small targets [2].

3. CLUTTER REDUCTION

In general, detection is a trade-off between detection probability and false alarm rate. Often, automatic detection algorithms for small targets at sea are based on statistical assumption about contrast of a target with background. Setting a threshold on contrast related to these statistical properties (for example the standard deviation for an assumed Gaussian distribution) allows some control of the amount of false alarms. When the threshold is set too high, small targets will not always be detected. Clutter reduction is a process of discriminating detections related to possible targets of interest from detections due to the background, allowing the false alarms to be filtered out. This allows a lower detection threshold (and higher detection probability), with a lower false alarm rate.

Since clutter reduction is part of the continuous real-time processing loop (see Figure 1), only simple features can be computed for the (possibly many) detections. Often, the size of a detection in pixels is used. However, this does not allow discriminating small targets at a distance, from sea clutter nearby. In an earlier project “SURFER” [3], a large dataset was used to determine features for small targets (including zodiacs and kayak), and sea clutter. It was found that simple features comparing the variation of intensity on a detection with that of nearby background, allowed for a good separation between targets and clutter for very different sea states and several bands.

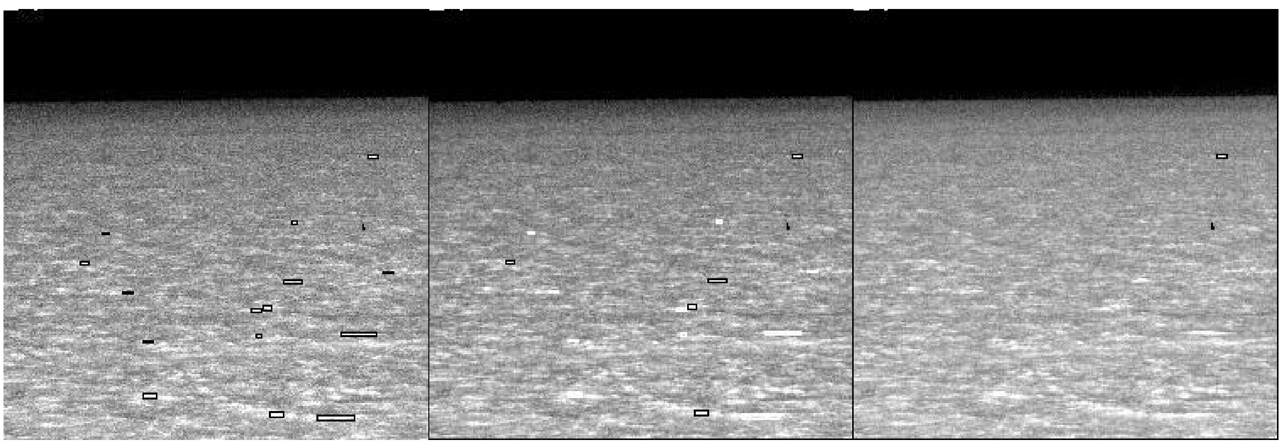

Fig. 4. Results of track filtering and clutter reduction. The left image shows all detections. The middle image shows only tracks that last for at least 15 frames (0.3s). The right image shows the result after filtering using features based on simple statistics of spatial target intensity. The only detection left is a kayak. Some bad pixels (black) are visible below the kayak.

Elimination of false alarms may also be done using tracking. Many causes of sea clutter are short lived, with only some white caps lasting several seconds. Filtering tracks that only last less than a fraction of a second will already remove many false tracks, as shown in Figure 4. In this mid-wave infrared recording of a kayak at sea, there are many detections due to waves, but filtering on track lifetime leaves only some white caps. These, and almost all other short-duration detections, are easily filtered out with the described features, as shown in the right figure. The thresholds used on these features were such that each would not filter all clutter, but the combination only rarely passes a white cap, and almost always detects the kayak.

4. ASPECT ANGLE DEPENDENCY IN MEASURED AND SIMULATED SHIP PROFILES

In operational use, when discriminating between classes, the total of classes should encompass all possible ships that are encountered. This may include an ‘unknown’ class for ships that are not described by the known features. Describing the classes by features, requires enough data to determine whether an encountered ship belongs to that class. For example, when classes are as detailed as types of frigates, data of all types should be available, and all ships for which there is no data, will be unknown. This requires large databases, that will most often not be complete. Further, many features depend on distance, aspect angle and even weather. Even simple features such as area or width of a detection are different when a target is seen from another distance or angle. Distance dependency may be eliminated if the distance (and therefore scale) is known. Some features, such as width/height ratio, do not depend on distance, but many discriminating ones, such as area, are. Dependence on aspect angle is more complicated.

When the aspect angle is not taken into account when determining the range for a feature for different classes, these ranges will often overlap. This prevents the use of this feature for separating these classes. Either features should be made invariant to aspect angle, or the aspect angle of the target should be known and the ranges for the feature should be known for different aspect angles. This can be helped by data from other sensors. For example, a radar track can give a distance and aspect angle estimate. A laser range profile provides complimentary information as it gives information perpendicular to that of the camera. Using a laser range profile for discriminating ships is shown in [7].

Making a silhouette aspect angle independent is not completely possible, as it depends on both width and length of the object, and parts may be occluded at certain angles. This means not only data is needed for many ships, but also for different aspect angles. This is especially true for more complicated features, such as profiles, used to correlate a measured silhouette to a database.

We looked at the influence of aspect angle on the profile in real measurements of a frigate (dataset 4 in Table 1) and on profiles determined from images simulated using a 3D model of the same type of frigate. The data has some limitations, because of segmentation artifacts in the measured data, the 3D model not being entirely correct, and not having an absolute measure how well profiles are alike, compared to those of other ships. It does however give a good indication of how aspect angle influences profiles.

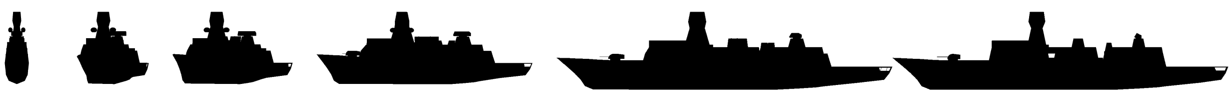

Fig. 5. Silhouettes simulated from a 3D model of a frigate, at aspect angles of 0, 10, 20, 40, 70 and 90 degrees.

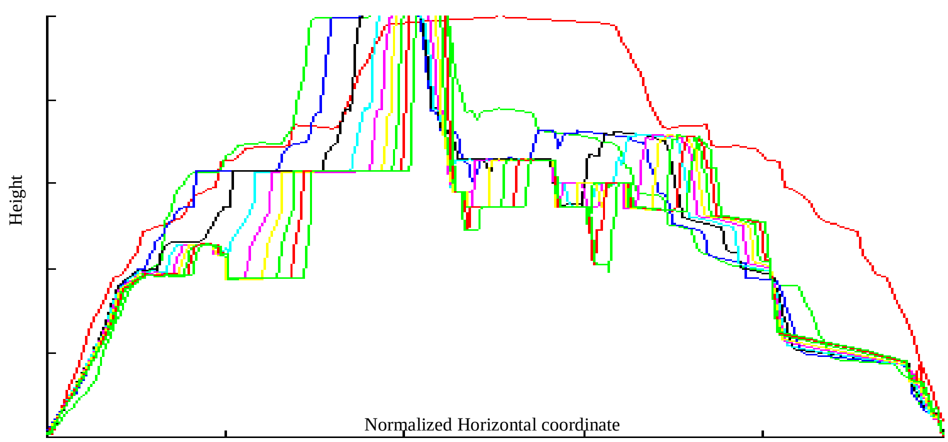

Fig. 6. Height profiles of a frigate simulated from a 3D model for aspect angles from 0 to 90 degrees. Width and height are normalized. Small aspect angles result in the largest profiles, with least intricate shape. Each line shows what the top of the silhouette looks like when viewed from a certain angle, stretched to a single width.

Figure 5 shows examples of simulated silhouettes of a frigate. It is clear that near head-on the appearance is very different from the side view. However, for angles in between, there are much less differences. This is also visible in Figure 6 where the profiles are shown for profiles of a simulated frigate at angles from 0 (head-on) to 90 (side view) degrees. The profiles are scaled to have the same width, and height.

The profiles show that there is only limited variation between many aspect angles, but when approaching a head-on view, the profile changes dramatically. The same is seen when looking at profiles from the real measurements.

To have a closer look at the similarities between height profiles, a difference measure was used, defined by:

$$D_{ij} = \sum_{x} \left| f_i(x) - f_j(x) \right| \tag{1}$$

Where indices indicate profile functions f for different angles. The difference is summed for all horizontal locations and is a measure for the area between the plots of two profiles. For this, all profiles are resampled to have the same number of points as the widest profile.

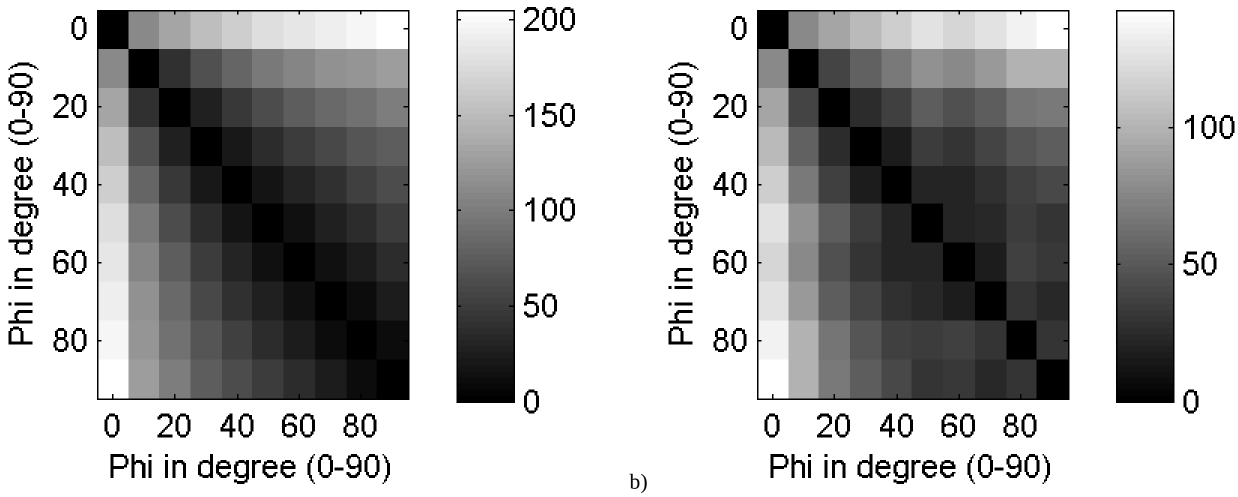

a)

Fig. 7. Difference between profiles of a frigate at different angles. a) for simulated profiles b) for measured profiles. A higher value (whiter in the figure) indicates a larger difference.

Figure 7 shows these values computed for all combinations between the profiles, both for the simulated and measured ones. Even though the values do not give an absolute measure, it is clear that differences between profiles for nearby angles (near the diagonal) are small. Further, the near head-on profiles are much different from all other profiles. Compared to those differences, profiles for angles from 20 to 90 degrees are very comparable, as seen in the bottom rows of the figures. Although for measured profiles the figure is a bit more irregular, the same conclusions can clearly be drawn for simulated and measured data.

5. DISCRIMINATION BETWEEN SMALL TARGETS

5.1 Introduction

In this section results are shown of ongoing research on the usability of features for discriminating between targets at different levels. Topics of this research are how well ships of different shapes can be distinguished, and what general information about size, shape or type can be obtained from simple features, even when the target is only observed as a few pixels high. The latter occurs for example when targets are at larger distance, and a system more aimed at detection, with corresponding smaller instantaneous field of view, provides the image of the target. An example of such a system is the sensor described in [8], which is a staring array, consisting of uncooled long-wave infrared cameras for detection, and color cameras for observing the targets. As 12 cameras have to span 360 degrees, resolution on a distant (small) target is limited.

Features that were examined are simple ones (e.g., based on dimensions, intensity variation) and moments (describing shape). Both are influenced by target distance and aspect angle. In the next sections, the use of these features for discriminating small targets is discussed, applied to data recorded in Rotterdam harbor (dataset 3 in Table 1).

5.2 Rotterdam data

Data was recorded in Rotterdam harbor in the Netherlands for a rhib, a water-taxi (a small cabin boat used in Rotterdam to transport people) and a jet-ski [8]. Recordings were made in two infrared bands, and visible light. Examples of the rhib and water-taxi recordings in the longwave IR band (recorded with an AIM QWIP camera) are shown in Figure 8.

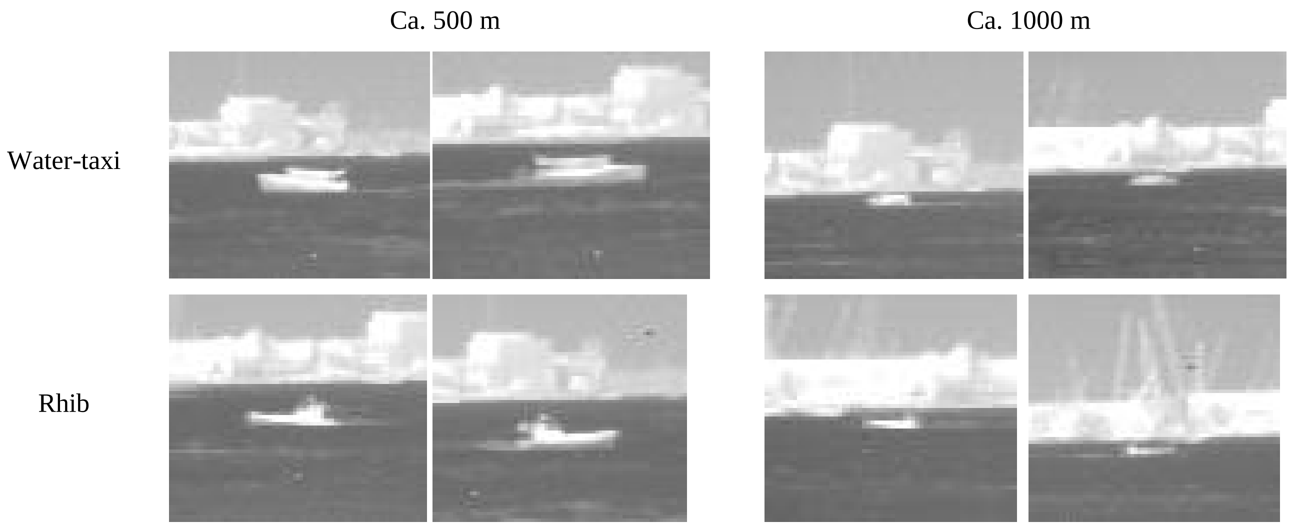

Fig. 8. Long-wave infrared recordings of water-taxi and rhib at different distances (indicated as planned distances of 500 and 1000m), in Rotterdam.

For distance, there is only a rough value for some recordings as a maneuver was planned at a certain distance. In order to examine the influence of distance and aspect angle, these values have to be estimated.

The distance is estimated from the camera height and the angle between the horizon and the bottom of the detection, where it touches the water. Since both camera height above the water and the horizon (automatically detected from the water-land transition) are not accurately known, the distance estimate is not perfect, but it does allow sorting detections by distance.

The aspect angle is estimated from the observed width of the detections, converted to real dimensions using the estimated distance and camera resolution. Taking the largest and smallest observed width, allows estimating the width and length of the targets. A simple formula describes how the observed width depends on these dimensions and aspect angle. The aspect angle can then be computed from the observed width of a detection by inverting this formula. Especially at larger distances, where the error in horizon location has the largest influence, the estimate is often incorrect. Due to estimating the horizon too low, the distance is overestimated, as is the width after correcting for distance. To compensate for this, the aspect angle (computed from this width) is scaled to better map between 0 and 90 degrees.

An algorithm for automatic detection of small targets at sea was used to obtain target segmentations of the boats, used in part of the results. This method is based on a constant false-alarm-rate detector and a background estimation that incorporates statistical changes in a sea background. This is an updated version of the detection method described in [3]. Although it works well in most cases, there were some difficulties in consistently detecting the targets when they were near the water-land border. As automatic detection is not part of this research, the targets were indicated manually, and an automatic segmentation was then performed at the indicated target position. In some cases, part of the wake of the boats may be included in the segmentation. No manual correction was made for this.

5.3 Simple features

A number of features are examined, based on:

- Dimensions of detection, in pixels. Most are scale independent (e.g., height/width ratio)
- Real dimensions, using the distance estimate (area, width, height)
- Intensity (total contrast, maximum contrast, standard deviation of intensity, standard deviation of intensity compared to background)

Most of these should be distance independent, either because they are not depending on scale, or because the distance estimate was used to convert pixel-based values to real dimensions. As an example, Figure 9 shows the estimated area as function of distance, for the three targets. The distance estimate becomes less correct at larger range, resulting in a general trend upwards. Still it is clear that in general, the jet-ski is distinguishable from the other two, and the rhib shows a slightly larger area than the water-taxi. It should be noted, that these data include all aspect angles, resulting in the large variations near distance values of 400 and 800 m, where the targets went around in a circle. At largest distances, detections are only a few pixels high, resulting in larger variations as well.

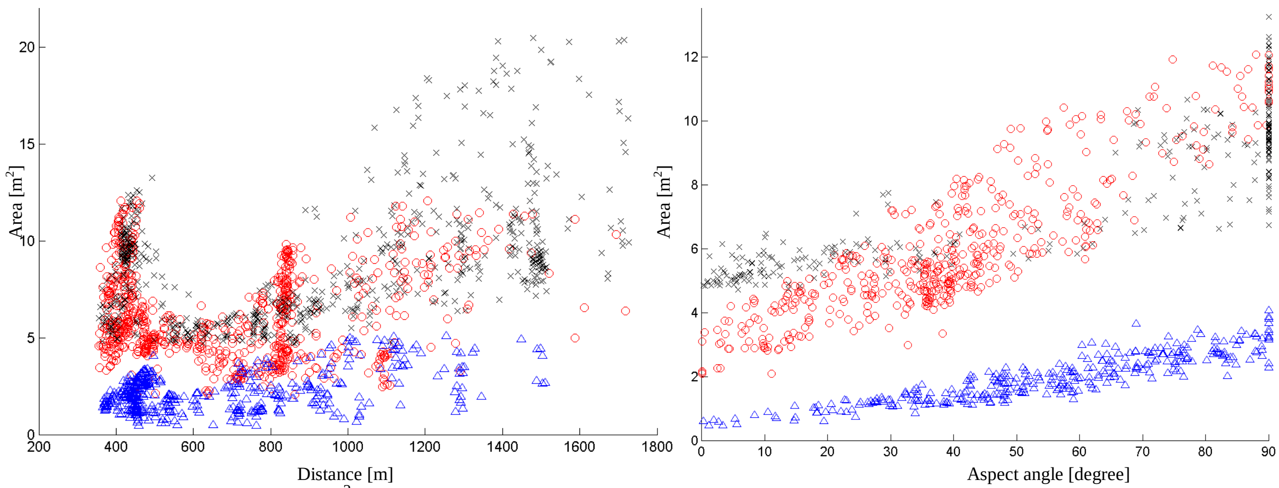

Fig. 9. Estimated area (in m^2), left as function of distance (in m), right as function of aspect angle (rough estimate, in degrees). Triangles: jet-ski, circles: water-taxi, crosses: rhib

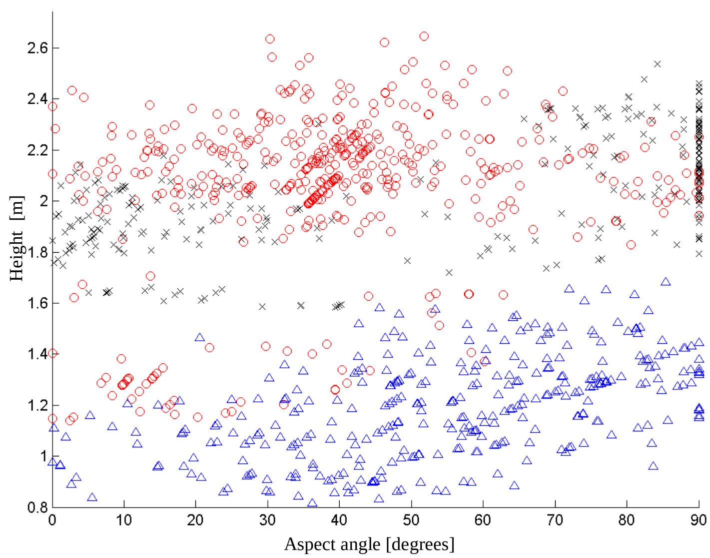

Fig. 10. Estimated Height (in m) as function of estimated aspect angle. Triangles: jet-ski, circles: water-taxi, crosses: rhib

When plotting the area as function of aspect angle, it becomes however clear that there are differences between values for rhib and water-taxi. Knowing the aspect angle would in this case help, and make a distinction between rhib and water-taxi possible in many, but not all, cases.

Area could also be corrected for aspect angle in much the same way as aspect angle was estimated from width as explained above. This is the case for all features that depend on aspect angle because they are computed from the detection width. In our case, this would however result in a simple dependence defined by our formulas, and give no insight in the actual behavior. Some features are already aspect angle independent, as Figure 10 shows for estimated height.

Using several features will allow to make a distinction between the three targets in many cases. Note that all values shown in the figures are the many occurrences of the same targets in an image sequence. When the majority shows a clear separation of features for the different targets, combining observations will increase reliability of the information.

5.4 Moments

For classification of ships, often more complex features are used, that describe the distribution of the detected silhouette, using segmentations of observed ships consisting of many pixels. Examples are Fourier descriptors [11,14] and moments [1, 4, 10, 11]]. Moments are computed using the difference of the coordinates of each detected point to an average value, raised to a certain power and summed. Moments describe features as size ratios, symmetry, eccentricity, etc. An example of moments applied to some of the targets from visual light recordings (dataset 1 in Table 1), are shown in Figure 11.

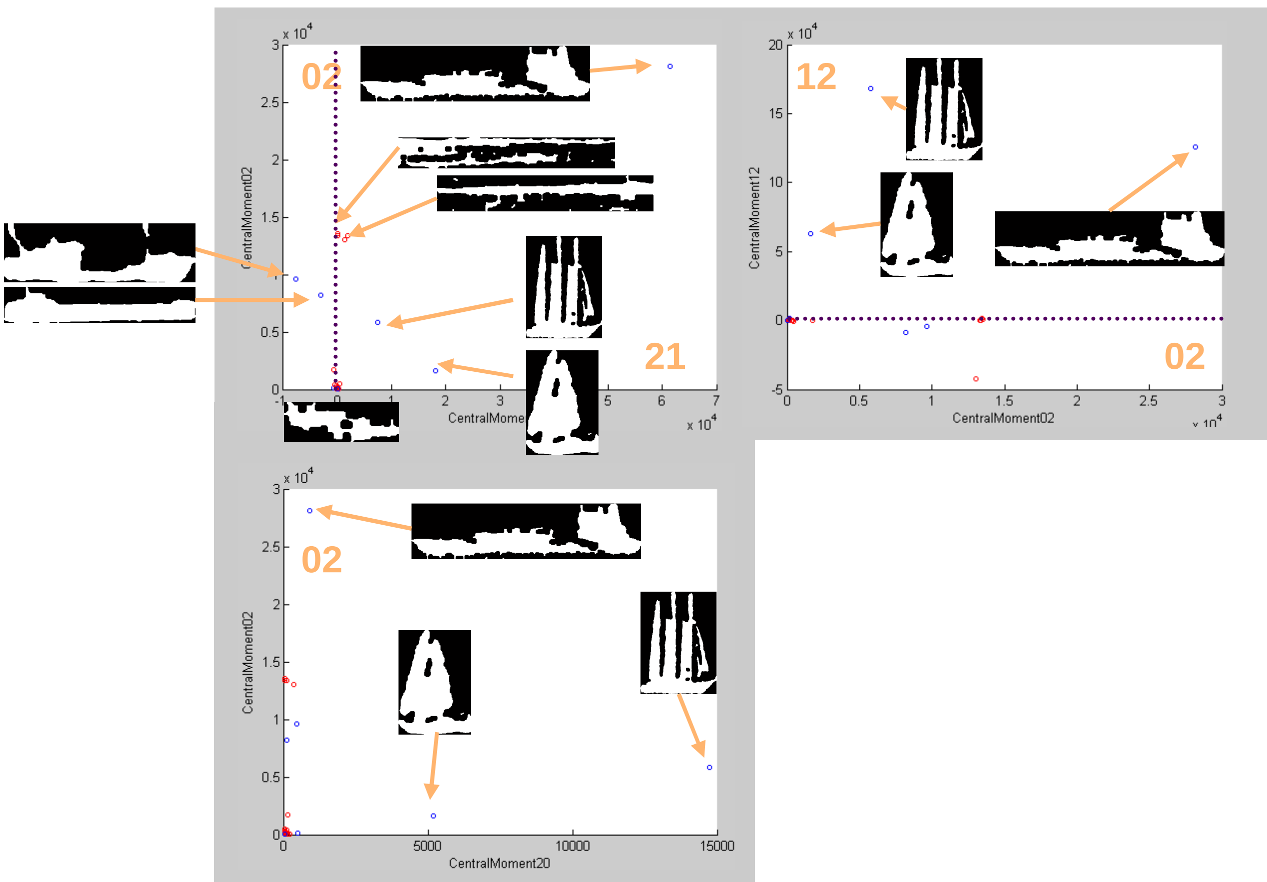

Fig. 11. Moments computed for different ships and some false alarms. Varying combinations of the shown moments allow differentiation between different ships, based on ratios and shape.

There, combinations of moments are plotted. Trends can be seen separating different shapes, such as flat (container ships), high (sailing boats) and irregular (clutter).

As moments are computed from pixel coordinates (see equation 2 for central moments), they depend on distance, and camera resolution:

$$\mu_{ij} = \sum_x \sum_y (x - x_c)^i (y - y_c)^j f(x, y) \tag{2}$$

where *i* and *j* indicate the order for *x* and *y* direction. (x_c, y_c) is the centre-of-mass, and function *f* in our case is equal to 1 or zero. A modified version allows computation of more scale invariant moments (equation 3 [13]), defined by:

$$\rho_{ij} = \frac{\mu_{ij}}{(\mu_{00})^{(i+j)/2+1}} \tag{3}$$

These moments, up to order 3 (resulting in 16 moments), were computed for the three targets. Figure 12 shows three of these moments. The middle picture shows the moment of order (0,3), which is related to horizontal pixel distribution, seen from the side at distances of 500m and 1000m, and going in two directions. The water-taxi has an even distribution of detected pixels, as its value is close to zero and there is no difference when going to the left or to the right. The rhib clearly has its pixels distributed unevenly, resulting in a non-zero value, and sign change when changing direction. The jet-ski is a bit more irregular. For the larger targets, there is no clear difference between the values at the two distances, showing that the computed moments are indeed scale invariant. When looking at the different moments shown in the figure, it can be seen that for some moments, the difference between two of the targets is small, and they may not be distinguishable. However when looking at other moments, these two targets may show large differences. For example, jet-ski and water-taxi have almost the same moment of order (0,2), but show clear differences for order (3,0).

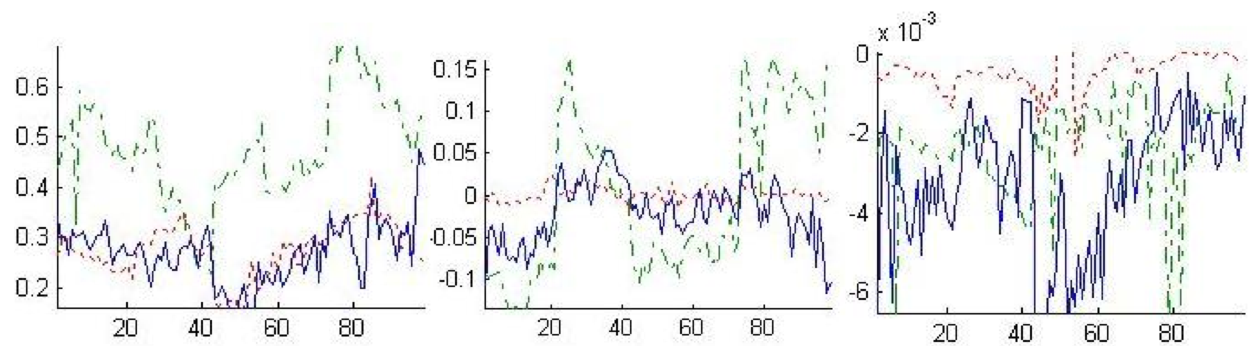

(0,2) (0,3) (3,0)

Fig. 12. Scale invariant moments for rhib (-.), water-taxi (--) and jet-ski (-), at distances of 500 (x values 1 to 40) and 1000m (x values 41 to100), for side views, going from left to right (x 1-20 and 41-70) and right to left (x 21-40 and 71-100). Orders (0,2),(0,3) and (3,0) are shown.

These results are for side views only. When looking at values for all aspect angles (estimated as described in Section 5.2), differences are less clear, as shown in Figure 13, where the moment of order (0,2) is shown as an example. Looking at a range of values that can occur for a target over all angles, it is clear that there is a huge overlap between the targets. Without any knowledge of the aspect angle, it will not be possible to distinguish the targets based on these moments. When an estimate is known however, there are clear differences, with much smaller overlap. The order (0,2) moment for example, shows that the water-taxi generally has the lowest value, rhib the highest, with jet-ski somewhere in the middle, but overlapping with the others. Combining different moments may then allow a discrimination between the targets. This is shown in Figure 14, where the (0,3) and (3,0) moments are plotted against each other, for the side views of the three targets. Although there is still some overlap, there are clear clusters for the three targets, with the rhib being split in two (for two different orientations). Several classifiers were used to automatically determine these clusters. The figure shows three, computed using the PRTOOLS software [5].

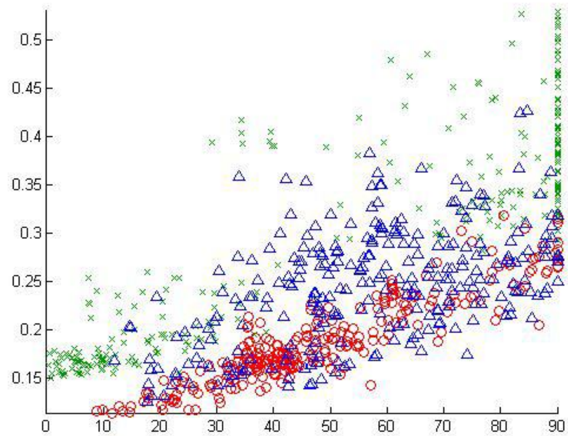

Fig. 13. Scale invariant moment order (0,2) for rhib (crosses), water-taxi (circle) and jet-ski (triangle) as function of estimated aspect angle.

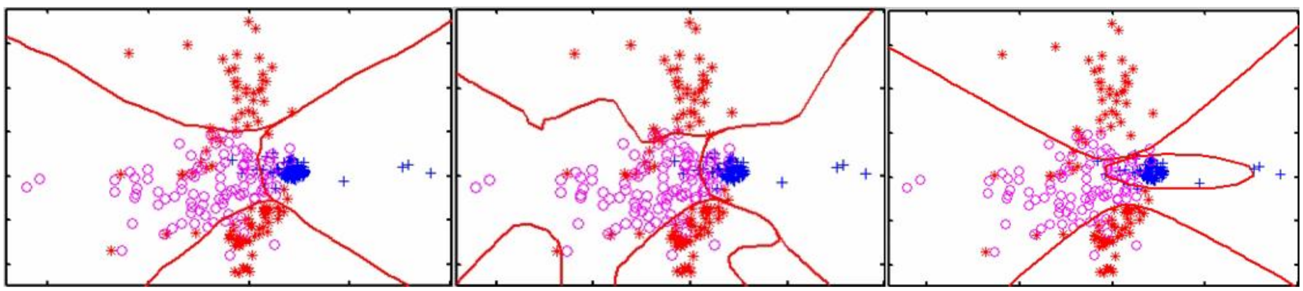
Fig. 14. Results of different classifying methods separating clusters for rhib (crosses), water-taxi (plus-signs) and jet-ski (circle), using k-nearest neighbor classifier, Parzen classifier, Normal densities based quadratic classifier.

6. CONCLUSION

This paper discussed the procedure of adding information to detected objects, by using features to label a target as belonging to different classes. Results from ongoing research were shown that illustrate aspects of this process, and that are of importance in labeling (small) targets in a coastal environment. Clutter reduction is an example of very basic labeling. It is shown that false alarms due to clutter can effectively be removed using track information and by using simple features such as statistics of target intensity variations.

Height profiles of a measured and simulated frigate were examined for influence of aspect angle. It was observed that most variation is in a limited range of angles near the head-on view. Variation from about 30 to 90 degrees was limited, when normalizing the profile for observed width.

The presented results obtained using simple features and moments for small targets, show that it is possible to use them to discriminate between small targets, even when their segmentation is only a few pixels high. A single one of the features may not be enough, but a combination of features allows to distinguish between targets in many cases. Using several scale invariant moments, a clear difference is visible, but information about size is lost, and intensity of the detection is not taken into account. It makes therefore sense to combine such features with simple features such as area, and intensity variations over the target. The latter could also be used to compute more complicated moments.

Features should preferably be invariant to how the object is viewed, such as scale (due to distance and camera resolution) and aspect angle. Where distance dependencies can partly be corrected for and introduce mostly inaccuracies, different aspect angles may result in entirely different values. It is clear that in order to be able to distinguish between targets, having an estimate for the aspect angle is a necessity. Such an estimate may be obtained from track information from other sensors (such as radar), or roughly estimated from the data, if enough observations at different angles are made.

Further work is being done on the described and other features, clutter reduction, labeling and visualization in order to assist an operator in identifying targets of interest.

7. ACKNOWLEDGEMENTS

The work for this paper was supported by the Netherlands MoD under the programme V602 "Electro-Optical sensor systems". We would like to thank FFI in Norway for providing the detailed data of the frigate recorded during the MCG/8 NATO trial in Stavanger in 2007.

REFERENCES

[1] Alves, J., A., "Recognition of Ship Types from an Infrared Image Using Moment Invariants and Neural Networks," Masters Thesis, NAVAL Postgraduate school, Monterey (2001).

[2] Broek, A.C. van den, Broek, S.P. van den, Heuvel, J.C. van den, Schwering, P.B.W., Heijningen, A.W.P., "A multi-sensor scenario for coastal surveillance," SPIE Proc. 6736-24, 67360P6 (2007)

[3] Broek, S.P. van den, Bakker, E.J., Lange, D.J.J. de and Theil, A., "Detection and classification of infrared decoys and small targets in a sea background," SPIE Proc. 4029, 70-80 (2000)

[4] Dudani, S., Breeding, K. and McGhee, R., "Aircraft identification by moment invariants," IEEE Trans. On Computers, 26, No.1, 39-46 (1977)

[5] Duin, R.P.W., Juszczak, P., Paclik, P., Pekalska, E., de Ridder, D. and Tax, D.M.J., "PRTools4, A Matlab Toolbox for Pattern Recognition," http://prtools.org, Delft University of Technology, 2004

[6] Gouaillier, V. and Gagnon, L., "Ship Silhouette Recognition Using Principal Components Analysis," SPIE Proc. 3164, 59-69 (1997)

[7] Heuvel, J.C. van den, Pace, P., Bekman, H.P.Th., Putten, J.M. van, and Schleijpen, H.M.A., "Experimental Validation of Ship identification with a Laser Range Profiler," SPIE Proc. 6950-29, (2008)

[8] Kemp, R.A.W, de Groot, J.F., Broek, S.P. van den, Lange, D.J.J. de, Dijk, J. and Schwering, P.B.W., "Results of Optical Detection Trials in Harbour Environment", SPIE Proc. 6943-34, (2008)

[9] 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 (2004).

[10] Lauberts, A., Karlsson, M., Näsström, F. and Aljasmi, R., "Ground Target Classification Using Combined Radar and IR with Simulated Data," Proc. of the 7th international Conference on Information Fusion, vol. 2, 1088-1095 (1998)

[11] Reeves, A., Prokop, R., Andrews, S. and Kuhl, F., "Three-dimensional shape analysis using moments and Fourier descriptors," IEEE Trans. on PAMI, 10, 6, 937-943 (1988)

[12] Schwering, P.B.W, Broek, S.P. van den and Iersel, M. van, "EO System Concepts in the Littoral", SPIE Proc. 6542-100, 654230 (2007)

[13] Sonka, M., Hlavac, V. and Boyle, R., "Image Processing, Analysis, and Machine Vision," Brooks/Cole Publishing Company, 1999

[14] Wallace, T. P. and Wintz, P. A., "An efficient three-dimensional aircraft recognition algorithm using normalized Fourier descriptors," Comput. Graphics Image Processing 13, 96-126 (1980)