Inference of vessel intent and behaviour for maritime security operations

Bert van den Broek [*], Arthur Smith, Eric den Breejen, Imelda van de Voorde

TNO, Oude Waalsdorperweg 63, 2597 AK The Hague, The Netherlands

ABSTRACT

Coastguard and Navy assets are increasingly involved in Maritime Security Operations (MSO) for countering piracy, weapons and drugs smuggling, terrorism and illegal trafficking. Persistent tracking of vessels in interrupted time series over long distances and the modelling of intent and behaviour from multiple data sources are key enablers for Situation Assessment in MSO. Results of situation assessment are presented for AIS/VTS observations in the Dutch North Sea and for simulated scenarios in the Gulf of Oman.

Keywords: Maritime situation awareness, situation assessment, information fusion, anomalous behaviour, vessel intent.

1. INTRODUCTION

Due to global economic and socio-political changes, an increase of conflicts near the world's coastlines is anticipated. The current conduct of expeditionary missions, Maritime Security Operations (MSO), and Peace Support Operations (PSO) means that Navies have to control instead of dominate the sea. This implies allowing regular vessel traffic in the area of operations, and act against irregular adversaries who nevertheless also can possess military armaments. In this combined military/non-military setting of operations, acts of piracy, drug smuggling and other threatening events become obscured in the crowd of everyday fisheries, cargo traders, ferries and pleasure crafts.

Maritime Security Operations are conducted from an international military as well as from a border security and Coastguard perspective. The military and security requirements are much the same in the sense that a detection capability for small vessels and anomalous behaviour of vessels, large and small, reporting and non-reporting, are needed. The difference between the border security and the military out-of-area operations is that the area to be covered is a few orders of magnitude larger for the latter.

In the home-based Coastguard perspective drugs smuggling, illegal fishing and illegal trafficking are adverse intents that need to be recognised. This is not only a national responsibility and need to be addressed by the countries in the region. For example, Europe has embedded most of its research and development of surveillance and information systems for border control in the EUROSUR programme. Increasing the level of shared situational awareness at the external borders and improving the reaction capabilities of national authorities surveying the borders is essential here.

Current military operations are missions to act against piracy, illegal fishing, smuggling of people, and terrorism. Examples of ongoing military missions are Active Endeavour in the Mediterranean Sea, and Ocean Shield in the Horn of Africa. For these out-of-area operations, surveillance implies that all vessels within the designated area should be tracked, identified, analysed and considered for possible boarding, using all available sensors. Every contact determined to be a suspect vessel should be tracked persistently and considered to be a potential target for a boarding operation.

In this paper we present an overview and results of our Maritime Situational Awareness research programme on semi-automatic picture building for the home-based Coastguard as well as out-of-area maritime security operations. This programme covers topics such as persistent tracking, information fusion for anomaly and intent determination as well as decision support for asset planning. Based on the results we have developed a test bed for the home-based Coastguard context aiming at real time processing and anomaly detection using data from sensors such as coastal radar and AIS data as part of a vessel traffic system (VTS). For the out-of-area context we developed a simulation test bed for a tactical

[*] bert.vandenbroek@tno.nl; phone +31 888 66 4075; http://www.tno.nl

Unmanned/Unattended Sensors and Sensor Networks X, edited by Edward M. Carapezza, Panos G. Datskos, Christos Tsamis, Proc. of SPIE Vol. 9248, 92480E · © 2014 SPIE
CCC code: 0277-786X/14/$18 · doi: 10.1117/12.2073676

decision aid that allows optimizing the maritime picture by automatically recognising intents using sensors on-board mobile surveillance assets, such as frigates, UAVs and helicopters.

This paper is organised as follows: in section 2 we discuss the properties of Information Guided Operations (IGO) being the core item for MSO. Then in section 3 we describe the methodology for intent recognition. In section 4 and 5 we elaborate on the test beds for the home-based Coast guard and the out-of-area context, respectively, before we finish with conclusions in section 6.

2. INFORMATION GUIDED OPERATIONS

MSO are increasingly becoming Information Guided Operations (IGO). Required item for IGO is an effective operational picture of vessel behaviour and vessel intent to ensure that further ships inspections are successfully revealing adverse ship intents as much as possible. The approach in information operations can be described in two steps.

Step 1: determine whether a vessel belongs to a normal category or not. In case it is not normal, further inspection strategies are to be considered. For this early warning step information is collected by sensors of the vessel traffic service (VTS) or by observation at distance from monitoring assets on board of a frigate or airborne assets such as UAVs and helicopters.

Step 2: this step comprises further close-in inspection by surveillance assets, providing extra information, which eventually may result in boarding. Since further surveillance requires extra effort, may warn the possible opponent, and may result in boarding of an innocent vessel a cost-benefit analysis needs to be made by assessing the possible adverse intent and its impact. This analysis is directed by the operation were also rules of engagement etc. are specified.

In the home-based Coastguard context normal traffic is predominant and hostile intent (e.g. smuggling) is only a small fraction of all vessels of which the impact on the security situation is also limited in general. The first step (normal/anomalous separation) is sufficient to optimise the further close-in inspection strategy.

In the out-of-area context the normal behaviour is often less defined and known. Also the hostile intent and its impact on the security situation is often significant (piracy case) and the reason of starting the security operation. It is therefore the main topic of an out-of-area operation to assess an intent as early as possible. In the table below we summarise both steps versus the context.

Table 1: relevance of approaches in the operational context.

	Step 1 Normal-anomalous separation	Step 2 Adverse intent assessment
Home-based Coastguard context	++	+
Out-of-area context	+	++

Nevertheless is understanding of the normal traffic important in out-of-area operations, since it will provide information for IGO which will reduce the number of unnecessary inspections and boardings. This is one the reasons of the *Dhow Project* (http://www.shipping.nato.int/Pages/Dhow-Project.aspx) started as a combined effort of the counter piracy operations. However this takes time and the importance of step 1 may become more important in the aftermath of the operation when the security situation should be evolved to more normal.

Various studies have focused on anomaly detection [1,2] and on close-in adverse intent recognition in case of out-of-area operations, where obvious intent indicators are present (e.g. vessel that approaches high value target at high speed, [3]). In the following we focus on the early warning case in which no obvious indicators are present and that less discriminative indicators have to be combined to reveal the hostile intent.

3. METHODOLOGY FOR ASSESSING ADVERSE INTENTS

To reveal intents we look for observable indicators that can detect signatures of the modus operandi related to the vessel intent. Ideally this indicator is specific for the intent and always present so that observation of the indicators directly implies detection of the intent. For example a go-fast crossing the Caribbean Sea is a good indication of drug smuggling, since other traffic is not showing such behaviour.

Most ships with adverse intent however try to hide within the daily normal traffic so that such ideal indicators are normally not present. A single information source for determining intent is therefore out of the question. We have to look for a combination of less specific indicators. Three questions arise here:

1. how to find useful indicators
2. how to assess indicators
3. how to combine indicators

To answer the first two questions ideally information is collected and analysed with respect to normal and abnormal behaviour. In the home-based Coastguard context this information can be collected with Vessel Traffic Systems (VTS). The collection of ground truth data for out-of-area operations is more difficult. For example piracy takes place in areas where often local commercial activities are done with non-reporting ships that do not carry AIS (such as Dhows).

Adverse events generally occur with very low probabilities. Moreover, the modus operandi for each type of adverse intent may change over time, which implies that relevant parts of the data remain scarce even after elapsed time. It is therefore difficult, if not impossible, to utilize the same methods as in home-based operations to ascertain a ground-truth of normal behaviour and to determine which kinds of data indicate anomalous behaviour. A possibility is to use estimations made by experts on the specific operational context. These experts have to propose indicators and to do the assessment.

3.1 Finding indicators

Since information in the database about adverse intent is scarce experts have to use their experience and imagination to propose indicators. To be useful for revealing the intent the different indicators need to be independent as much as possible. Since indicators are often quite qualitatively described it is not easy to proof this. It is therefore helpful to choose indicators from different viewpoints. We propose here to choose sets of indicators that belong to the following four categories, since these categories provide independent views. The categories comprise information about the objects itself, its behaviour, its location and its antecedents. We give here two examples for each category.

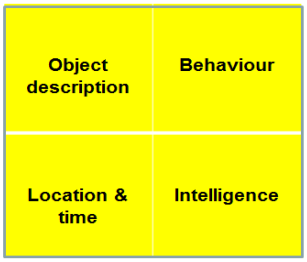

Object description
1) Dhow
2) Climbing device on board

Behaviour
1) Loitering
2) Heading change

Location & time
1) In fishery area
2) In PAG[1] area

[1] pirate action group

Intelligence
1) Malicious owner
2) Not registered

Figure 1: Left: the four categories shown in quadrants for finding indepedent indicators. Right: exampels of indicators per category.

3.2 Assessment of indicators

Once an indicator is proposed it needs to be assessed for its significance to discriminate the intent. Therefore its exclusiveness and presence w.r.t. the intent and other ships needs to be assessed. This implies to assess the parameters β (fraction of intent vessels that do comply with indicator) and γ (fraction of non-intent vessels that do not comply with indicator) shown in the figure below.

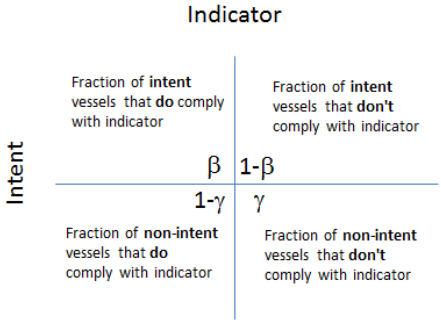

Figure 2. β and γ parameters which need to be assessed to determine the usefulness of an indicator to discriminate an intent.

Ideally β=1 and γ=1 (or β=0 and γ=0). In that case a single indicator suffices to detect an intent, like the go fast described in the introduction. However in practice indicators are not always true for intent vessels → β<1 and indicators are usually not exclusive for intent vessels → γ < 1. This means we have to use sets of non-ideal indicators and the question arises how to optimally combine these indicators to assess the possibility for detecting the intent given the set of indicators.

3.3 Information fusion and intent assessment

For combining the indicators we propose here 2 methods:

1. use of rules → RBS (Rule Based System)
2. use of the Bayes algorithm → BBN (Bayesian Belief Network)

for evaluating hypotheses of intent versus non-intent. The different kinds of intent (piracy, illegal fishery, etc.) are considered here to be independent.

A rule based system has the advantage that it is transparent to the user and that domain expertise can easily be used for assessing the rules. The Bayesian classifier is mathematically more refined but results are less easy to interpreted and accurate assessment of indicators are needed to take advantage of its refinedness.

Rule Based Systems (RBS)

The use of rules is common practice by operators. We use a version where the rules are tuned with weights (W_{pos}) that support a positive identification of the intent and weights (W_{neg}) that support denial of the intent. These weights can be derived from β and γ following:

$$W_{pos} = \frac{\alpha\beta}{\alpha\beta+(1-\alpha)(1-\gamma)} \tag{1}$$

and

$$W_{neg} = \frac{(1-\alpha)\gamma}{\alpha(1-\beta)+(1-\alpha)\gamma} \tag{2}$$

where α is the fraction of intent ships in the total population of ships considered. To get an assessment of our belief *B* in the intent we use the following separation function:

$$B(ind) = \frac{1}{(1+e^{a(b-ind)})} \tag{3}$$

where *a* and *b* are tuning parameters and *ind* is an index parameter based on the weights which only contribute when a rule is true (i.e. triggered):

$$ind = \frac{\sum_{tr} W_{pos} - \delta \sum_{tr} W_{neg}}{\sum_{tr} W_{pos} + \delta \sum_{tr} W_{neg}} \tag{4}$$

where δ is a normalization constant based on all weights used in the system for revealing this specific intent:

$$\delta = \frac{\sum W_{pos}}{\sum W_{neg}} \tag{5}$$

Bayesian Belief Networks (BBN)

Bayesian belief networks are standard tools for working with uncertain information. We use a basic version, i.e. a naïve Bayesian network [4], where the intent is the hypothesis H_k and the single node layer is given by the indicators I_1 .. I_n.

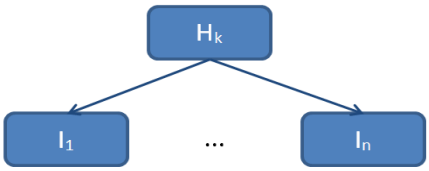

Figure 3: Depiction of the Bayesian network used.

We have to evaluate both the positive as well as negative hypothesis following

$$p(H_k|I_1..I_n) = p(H_k)\prod_{i=1}^{n} p(H_{k)}\, p(I_i|H_k) \tag{6}$$

$$p(\neg H_k|I_1..I_n) = p(\neg H_k)\prod_{i=1}^{n} p(\neg H_k)\, p(I_i|\neg H_k) \tag{7}$$

respectively and were

$$p(H_k) + p(\neg H_k) = 1 \tag{8}$$

Important here are the determination of CPTs (conditional probability tables) following

$$CPT = \begin{bmatrix} p(I_i|H_k) & 1-p(I_i|\neg H_k) \\ 1-p(I_i|H_k) & p(I_i|\neg H_k) \end{bmatrix} = \begin{bmatrix} \beta & 1-\gamma \\ 1-\beta & \gamma \end{bmatrix} \tag{9}$$

Furthermore, the global intent prior probability is given by

$$p(H_k) = \alpha \tag{10}$$

Note that the belief *B* for the RBS is comparable to the hypothesis probability $p(H_k|I_1..I_n)$.

3.4 Measure of effectiveness of the information fusion system for IGO

For evaluating the information fusion procedures described in the previous section we use here the following measure of effectiveness (*MOE*) for finding the intent ships:

$$MOE = POD - \kappa . FAR \tag{11}$$

where the *POD* is the probability of detection of the intent ships, i.e. the number of truly detected intent ships versus the actual number of intent vessels. *FAR* is a false alarm rate defined here by the number of falsely designed intent vessels over the actual number of intent vessels. κ is a cost-benefit factor. In IGO context detection is considered as a benefit, however finding a false alarm is spilled effort which is considered as cost. κ is a cost-benefit factor, which determines how much costs are allowed compared with respect to the benefit, and which is specified for the IGO.

Ideally the *MOE*=1, when the *POD*=1 and the *FAR*=0. The range for the *FAR* depends on the fraction of intent vessels α, so that for small α the *FAR* can in principle be much larger than the *POD* and negative *MOE*s are possible. κ of course is determined by the objective of the military operation. We will use κ=1 in the following.

3.5 Method assessment

Results are generated for various sets of α, β or γ. In general the RBS and BBN give comparable results for larger numbers of indicators (> 2) For smaller number of indicators the dispersion in the results for the RBS is larger compared to those for the BBN, so that the RBS is less reliable for only one or two indicators.

In the analysis we use here a so-called nominal good case with $\beta \in [0.8,0.9]$ and $\gamma \in [0.6,0.7]$, α=0.1 and a nominal bad case with $\beta \in [0.7,0.8]$ and $\gamma \in [0.4,0.5]$, α=0.1. Note that the potential for discrimination is high for both β and γ high (~1) , and for both β and γ low (~0). In the first case an indicator is called inclusive and in the second case called exclusive. Bad results are expected when values near 0.5 for β and γ are found. In the following figure we show *MOE* results using the BBN as a function of the number of indicators.

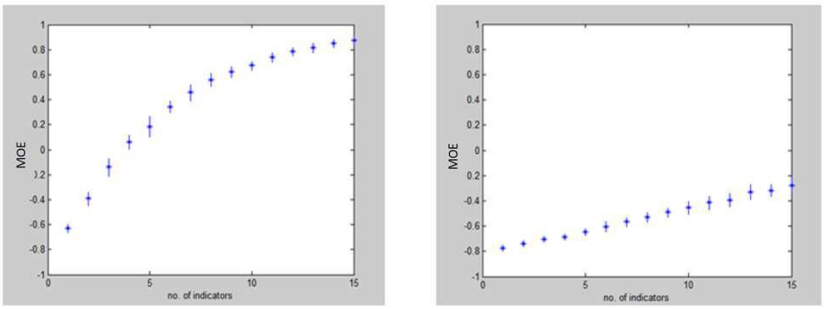

Figure 4: MOE as function of no. of indicators (α=0.1). Indicators have random numbers for β and γ in the range $\beta \in [0.8,0.9]$ and $\gamma \in [0.6,0.7]$ (left, nominal good case) and in the range $\beta \in [0.7,0.8]$ and $\gamma \in [0.4,0.5]$ (right nominal bad case). The bars indicate the spread in the results due to the various sets of random numbers.

The figure clearly shows that only reasonable results (*MOE* >0.5) are obtained in the nominal good case for 7 indicators or more. To explore the dependency of β and γ in more detail we use 8 indicators and have calculated the *MOE* as a function of β and γ (see following diagrams).

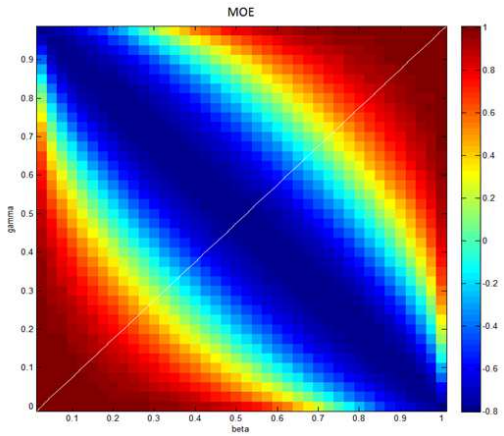

Figure 5: MOE as function of the weights β and γ (8 indicators, bins 0.04 for β and γ, α=0.1).

From the figure it clearly is found that a minimum is present for $\beta + \gamma = 1$, and maxima are found for $\beta = \gamma = 1$ or $\beta = \gamma = 0$. The behaviour of the *MOE* is also depending on α. In the figure below we show cross-cuts along the white line $\beta = \gamma$ in the previous figure for various value of α

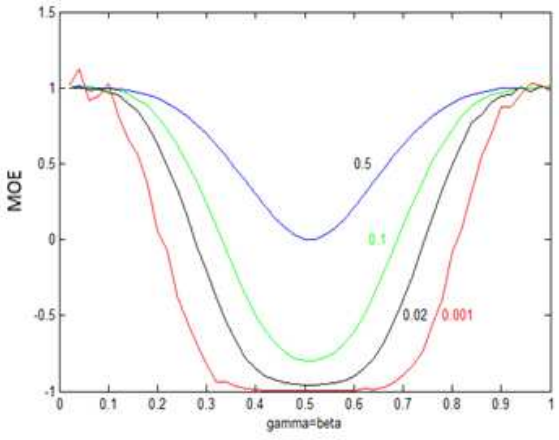

Figure 6: *MOE* as function of the line $\gamma = \beta$ for various values of α (see labels)

Clearly the criteria to find good indicators (both γ *and* β high or low) becomes more critical when α becomes smaller. The parameter $(\beta+\gamma)/2$ can be considered as a quality measure and below we give a table with the number of indicators needed to obtain *MOE* values > 0.5.

Table 2: indicator quality versus no. of indicators required

$(\beta+\gamma)/2$	No. of indicators
0.95	2
0.85	4
0.75	8
0.65	24
0.55	> 140
0.45	> 140
0.35	24
0.25	8
0.15	4
0.05	2

From this table we conclude that a reasonable and practical number of indicators implies that the following condition should apply at least: $(\beta+\gamma)/2)>0.75$ or $(\beta+\gamma)/2<0.25$. For other values an indicator does not substantially contribute and can better be removed. When proposing indicators it therefore important to do a good assessment of β and γ before the indicator is used. Also the assessment about the quality of the intent determination can then be made given the set of indicators by applying above described fusion methods.

Other methods

The weakness of the approach described above is that an indicator is first proposed and it should be tested afterwards to assess its usability for determining abnormal behaviour and vessel intent. There is an omission to find straightforwardly the most useful indicators at forehand. In a data mining approach, observational data from a group of vessels with normal behaviour is compared with data from a group of vessels with abnormal behaviour and/or adverse intent to find out which data representation can provide optimal discrimination. Of course such ground truth data need to be available.

4. REAL-TIME TEST BED FOR THE HOME-BASED COASTGUARD CONTEXT

For the home-based Coastguard context we focus on the actual data available from the vessel traffic services (VTS) in the English Channel and Southern part of the North Sea. The Dutch Coastguard made available a fused stream of coastal radar and AIS tracks from various posts. The live stream uses the open inter VTS exchange format (IVEF). We have built a real time test bed which uses this data stream to extract useful information and which applies rules to determine anomalous vessel behaviour. In order to obtain useful rules for anomaly detection the normal picture of vessel behaviour in the English Channel and Southern part of the North Sea has to be established. In the next section we present an analysis of open source data for vessel tracks for producing such a normal picture followed be a description of the test bed.

4.1 Constructing a normal picture

A structured information stream from VTS (AIS and radar) allows accumulation of a statistically rich database which can be analysed to produce a so-called normal density picture of ship traffic. Below we show results from an AIS data-set obtained through the Marine Traffic (MT) website (https://www.marinetraffic.com/) during one week of observation. Parameters extracted from the AIS data set are ship IDs, latitude and longitude, speed over ground (SOG) and course over ground (COG). The MT data-set show omissions since not all AIS data is publically received and not all recorded data will pass the internet for what kinds of reason. The data is therefore checked and missing data is repaired by interpolation if possible. Using spatial cells with dimensions of 1 by 1 km and using SOG> 0.05 m/s (no anchoring ships) number densities were generated, where the speed azimuth vectors were categorised using intervals of 10 degrees. In the figure below we show the results where the colour indicates the azimuth interval.

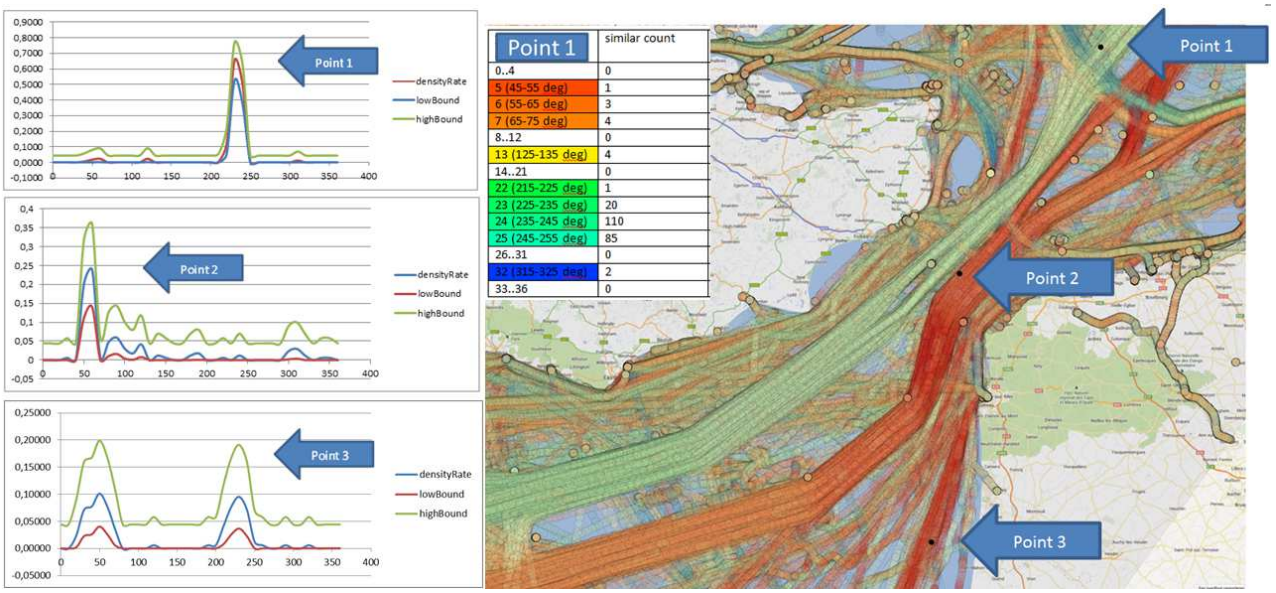

Figure 7: to the right density images of the English Channel with colours indicating the azimuth direction. For three points the azimuth distribution is also given (left)

4.2 Real time anomalous behaviour determination using complex event processing

In case of a continuous stream of monitoring data complex event processing (CEP) refers to events or patterns extracted in real time from the data. Applying appropriate combination and classification methods to the detected events provides instantaneous alarms in case of abnormal behaviour or adverse intent. We have developed such a CEP scheme based for real-time information extraction, called Scepter, and a real time combination method for detecting the maritime anomaly, called SeaBEAT.

In Scepter the event processing is performed in a hierarchical fashion, where a combination of low level events can trigger a higher level event. By intelligently filtering of the events, the data stream can be reduced to a manageable set. Scepter includes features to enable geospatial processing to be able to reason about the location and possible interactions of vessels. Important features Scepter supports are: 1) efficient processing of large quantities of data, 2) easy and quick adding of new events and patterns, 3) automatic reasoning about which underlying rules and events need to be activated to support a high-level event, 4) support for geographical processing, and 5) a web interface providing up-to-date event information. Examples of low-level behaviour events are ‘speed changes’, and ‘bearing changes’. Medium-level events can be ‘in zone’, ‘blacklisted vessel’, and ‘perpendicular to coast’. High-level events examples are ‘rendez vous in zone’, and ‘blacklisted vessel in zone’.

In an additional tooling (SeaBEAT) we have implemented a classification scheme to perform real-time behaviour analysis in the maritime domain. SeaBEAT uses the Scepter CEP framework for efficiently processing events, which can be split into four categories: information about the detection, location, trajectory, and behaviour events. SeaBEAT uses a rule based system for combining events and to determine the anomaly alarm. Both Scepter and SeaBEAT can be used as test bed and can be adapted to allow extraction of multiple and new events and to try-out more complex rule based systems.

Figure 8: SeaBEAT screen shot of NLD Maritime picture with alarms indicators after applying test rules for anticipated anomalous behaviour

5. TEST BED FOR THE OUT-OF-AREA CONTEXT

The programme of work for the out-of-area context aims at the development of a tactical decision aid (TDA) for optimising the maritime picture for large areas with a limited number of observation assets. Main topics are persistent tracking, information fusion and decision support for optimal deployment of the surveillance assets. To support the programme of work a suitable scenario has been defined and simulated using a recently developed simulation package (vessel traffic generator, VTG), [5]. Results are produced and visualised using the TDA test bed in which algorithms on basis of the simulation input can be tested.

Persistent tracking aims at fusing different, but distinct tracks of a vessel such that historical behavioural information becomes available. For example, a Dhow showing loitering tracks for several days in a certain area, is important information for evaluating its intent (e.g. fishery or piracy). The ship does not need to be continuously observed, but it is essential that the different tracks observed at various periods can be associated in the persistent tracking. For this purpose appropriate vessel recognition is important. Results of this study can be found in a related paper [6].

The information fusion aims at combining suitable indicators to predict the vessel intent. A large number of possible observable features have been evaluated as indicators for the intent following the procedure described in section 3.2. The more successful indicators in describing the intent have been combined using the rule based system, which has been described in section 3.3 and which is implemented in the TDA. Next to the RBS also other combination methods are considered such as a Bayesian network, a decision tree, a support vector machine and a neural network.

The intent predictions are used in the TDA to adjust and to optimise the initial planning of the surveillance assets, such as a frigate, UAV and helicopter carrying radar and electro-optical sensors, so that the improved observational information becomes available. In the TDA test bed several planning optimisation methods for surveillance assets deployment are studied [7].

5.1 Scenario and simulation description.

We have defined an out-of-area scenario in the Gulf of Oman and Arabic sea (see figure). This out-of-area scenario is the basis for the simulation of the data-set to be used in the test bed.

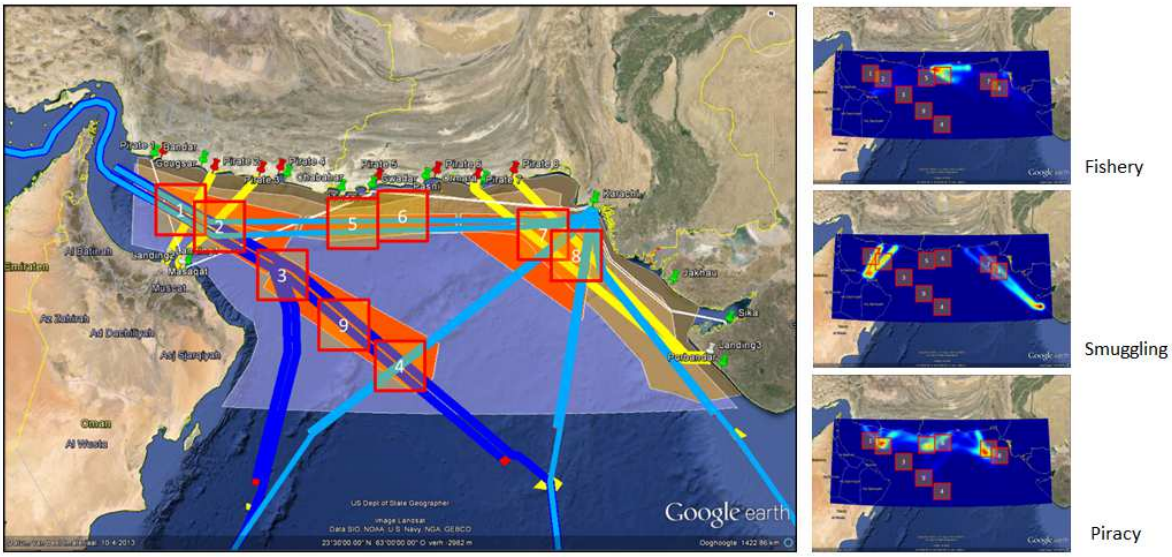

Figure 9: spatial layout of the operational with sea lanes, fishery (brown) and pirate action areas (orange), possible operational areas for a frigate with a helicopter and UAV are indicated (red squares). To the right the number density of simulated ships is depicted.

For the scenario three adverse intents (piracy, smuggling and illegal fisheries) were defined, next to regular and commercial activities (tankers, cargo ships, ferries and yachts). The scenario was populated by the VTG making use of the daily motion patterns of the different vessel types.

VTG in combination with the underlying VRforces (http://www.mak.com/pdfs/br_vrforces.pdf) provides high degree of interactions between the vessels including transition of states of the vessel activity. For example, a pirate vessel initially showing fishing behaviour switches to the activity of interception of a vulnerable commercial vessel when one is visually detected by the pirate and no Navy vessels are in the neighbourhood. The VTG automatically handles all the requested number of vessels for each vessel type, thus facilitating large vessel densities also of small non-reporting vessels.

All relevant attributes of the vessels, such as dimensions, flag state, cargo, people on board and maintenance state are statistically generated and dynamically updated by the VTG during scenario execution on the basis of its pre-determined intent. The attributes can be observed by the sensors of the surveillance assets, creating a perceived world. With the VTG generating all vessels in the entire simulated scenario the full set of ground truth is available for analysis and evaluation.

5.2 Tuning and use of the test bed.

Assuming ideal sensors that can detect all ships in the operational area, observation data were generated. From these data various possible indicators for the four categories in section 3.1 (behaviour, location and time, intelligence and ship properties) were tested by determining the weights β, γ (see section 3.2) and comparing these with the input intent type of the simulation. We present here results for the intent piracy. The simulation covered is about 1 day. In this simulation about 76000 ship detections (plots) were generated of which about 14500 have the piracy intent, so that α=0.2, relatively high to get reasonable test numbers for piracy. We found 8 indicators with the highest or lowest scores for $(\beta+\gamma)/2$ from the simulation results (see table).

Table 3: weights for selected indicators

Indicator type\Indicator weights	**β**	**γ**	**(β +γ)/2**
Mother vessel is present	0.35	0.98	0.67
Vessel has powerful propulsion engines	0.33	0.96	0.65
Climbing device is on-board	0.66	1.00	0.83
Weapons are on-board	0.70	1.00	0.85
Weapons are shown	0.49	1.00	0.75
Vessel is in pirate action area	0.60	0.83	0.72
Vessel owner is malicious	0.75	0.61	0.68
Vessel is registered	0.05	0.53	0.29

Most indicators are not so much positively indicating the piracy intent (moderate values for β), but very clearly indicating the absence of the indicator for the non-pirate ships (high γ). Note that the last indicator (lowest value) is a so-called exclusive indicator. i.e. being registered is a good indicator that a ship does not belong the piracy intent class, while most other vessels are (very low β and relatively high γ). We obtained piracy detections by applying the rule based system (see section 3.3) with above mentioned values for α, β, γ, and with a threshold of 50% for the piracy belief values. We found the following confusion matrix:

Table 4: confusion matrix for piracy vessels

	Detected piracy vessels	Detected other vessels
Actual piracy ships	0.18	0.01
Actual other ships	0.01	0.80

This confusion matrix shows that a very good separation (*MOE* = 0.9, see section 3.4) is obtained between the piracy and other ships using these indicators on basis of ideal observation data and that these indicators are suitable for use in the test bed. In using the TDA test bed the deployment planning of surveillance assets can now be tested using these indicators and weights, and using sensor models which take in account observation distances etc. This implies of course that detection results will be less optimal as in the ideal sensor case. By adjusting the deployment planning of the surveillance assets, the perceived maritime picture obtained through the observables collected with sensor data changes. The persistent tracking and the indicator extraction followed by information fusion discussed this section, can now be optimised. In the figure the display of the TDA test bed with a deployment of the observational assets is shown, where detection of ships is indicated together with a measure of effectiveness and other relevant information.

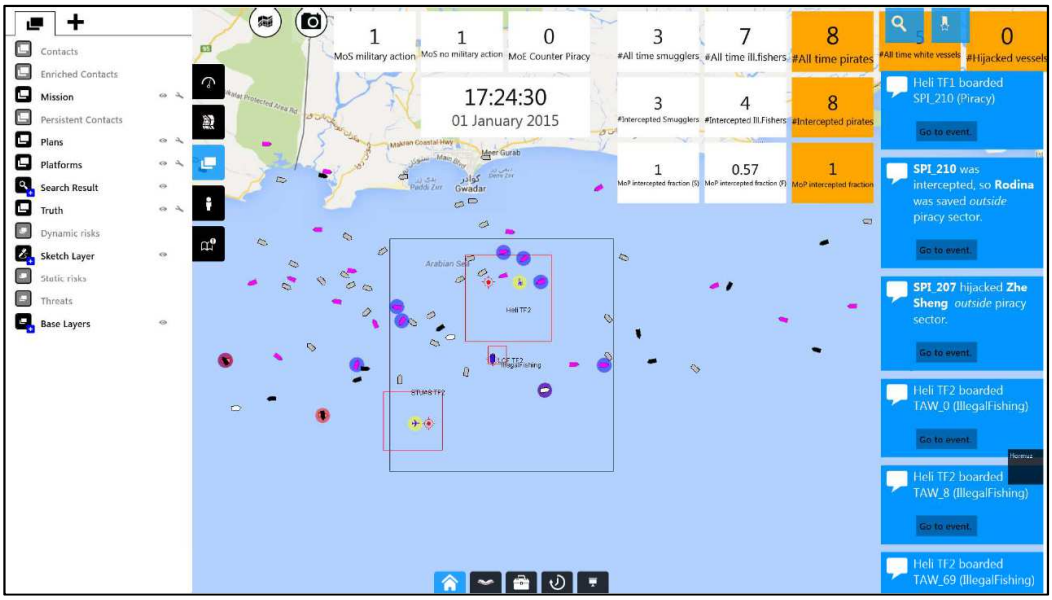

Figure 10: display of the TDA test bed.

6. SUMMARY AND CONCLUSIONS

We have discussed the concept of information guided operations (IGO) for maritime security in the home-based Coastguard context as well as in the out-of-area context and discussed the quality of indicators to distinguish adverse intents, such as smuggling, piracy, terrorism etc. We studied the fusion of indicators and found that for revealing an intent a limited set of meaningful indicators is clearly preferred over fusing a huge amount of easy accessible, but less discriminative features. As fusion method a tuned rule based system as well as Bayesian approach are adequate fusion methods.

For both the home-based Coastguard context as well as the out-of-area context we have presented a test bed for automatic picture building and situation assessment. The first test bed uses complex event processing and anomaly detection in a continuous stream of surveillance data from unattended sensors in a live sensor network. The second test bed, a tactical decision aid used in a simulated environment, aims at optimising the deployment planning of mobile naval assets. It is based on persistent tracking and generating an optimal maritime picture for situation assessment on basis of the information fusion results discussed in the paper.

ACKNOWLEDGEMENTS

The work for this paper was supported by the Netherlands MoD program V1114, Maritime Situational Awareness. Research leading to these results also has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 241598 (SeaBILLA project).

REFERENCES

[1] J. Roy, Automated Reasoning for Maritime Anomaly Detection, (2009), *NATO Workshop on Data Fusion and Anomaly Detection for Maritime Situational Awareness*, La Spezia, Italy, 15-17 September 2009

[2] Lane, R.O.; Nevell, D.A.; Hayward, S.D.; Beaney, T.W, (2010), *Maritime anomaly detection and threat assessment, 13th International Conference on Information Fusion*, Edinburgh, 2010.

[3] Bryan Auslander, Kalyan Moy Gupta, & David W. Aha, Maritime Threat Detection Using Probabilistic Graphical Models, (2012), *Proceedings of the Twenty-Fifth Florida Artificial Intelligence Research Society Conference*, May 23–25, 2012.

[4] Rish, I., (2001), An empirical study of the naive Bayes classifier, *IJCAI Workshop on Empirical Methods in Artificial Intelligence*, pp. 41–46.

[5] F Kuijper, Vessel Traffic Generator: Agent based maritime traffic generator, (2013), International Training & Education Conference (ITEC), 20-22 May 2013, Cologne, Germany

[6] S.P. van den Broek , H. Bouma, R.J.M. den Hollander, H.E.T. Veerman, K.W. Benoist, P.B.W. Schwering, (2014), Ship recognition for improved persistent tracking with descriptor localization and compact representations, *Proc. SPIE 9249-23*, (2014)

[7] Hilvert Fitski, (2013), Planning of Naval Assets for Maritime Situational Awareness, *Offshore Patrol & Security Conference, Portsmouth*, UK, 23 April 2013.