14th International Conference on Information Fusion
Chicago, Illinois, USA, July 5-8, 2011

Improving Maritime Situational Awareness by Fusing Sensor Information and Intelligence

A.C. van den Broek, R.M. Neef, P. Hanckmann, S.P. van Gosliga, D. van Halsema
TNO
The Hague, The Netherlands
bert.vandenbroek@tno.nl

Abstract - In present-day military security operations threats are more difficult to reveal than in conventional warfare theatres, since they take place during the course of normal life. These maritime missions often take place in littoral environments, where acts of piracy, drug trafficking and other threatening events become obscured in the crowd of everyday fisheries, cargo traders, ferries and pleasure cruises, hindering situation awareness. We aim to improve situation awareness and threat detection capabilities in maritime scenarios by combining sensor-based information with context information and intelligence from various sources. The fusion and analysis in order to reveal suspect from normal behavior is based on domain ontologies. A test bed allows the study of various exploitation and assessments techniques applied to these domain ontologies. Using an appropriate scenario we have simulated suspect and normal behaviour to test the applicability of the various techniques.

Keywords: Situation awareness, Intelligence, Sensors, Threat assessment, Decision support.

1 Introduction

Because of global economic and socio-political changes, an increase of conflicts near the world's coastlines is anticipated. The littoral zone is characterized by intense regular vessel traffic. The conduct of Maritime Security Operations and Peace support Operations means that navies have to control instead of dominate the sea, thus 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, naval forces have to protect themselves against threats from land, air and sea, while they continuously have to collect various data to ensure information superiority over their adversaries and thirds. For the purpose to achieve information superiority a research program at TNO, the Netherlands, has started aiming at improving maritime situation awareness. The study for this improvement focuses on the combined use of intelligence sources and sensor information. In this paper we first discuss the operational context and tasks, and define the information requirements. Next, we describe the fusion process, before we elaborate on fusion methods which can be applied. Finally we discuss a system architecture and simulation environment for testing the proposed methods and give conclusions.

2 Operational context

Security operations are often characterized by controlling large areas with a limited number of assets. An example is the anti-piracy operation in the Gulf of Aden where the operational area extends over thousands of sea miles. One of the main operational tasks is to direct assets timely to the right position. For the command and control process a common operational picture (COP) is the basis on which decisions and actions are taken. A core part of the COP is a maritime picture that contains information about vessel movements in the complex and detailed coastal environment and information about hostile intent and illegal activity of the vessels. Ideally the COP contains up to date information about the position of own, enemy and third entities combined with their missions, intentions, and capabilities.

In present-day military security operations threats are more difficult to reveal than in conventional warfare theatres, since they take place during the course of normal life. For example, during maritime missions in littoral environments, acts of piracy, drug trafficking and other threatening events become obscured in the crowd of everyday fisheries, cargo traders, ferries and pleasure cruises. The hostile intent of objects is therefore not always easy to determine because of its ability to cloak and hide among the regular vessel traffic.

To enable threat recognition appropriate situational awareness is needed which implies recognition of the objects present in the scene, their interaction with the environment and their intention on basis of threat hypotheses in order to foresee the situation in the near future. In this way hostile intentions and threats should be recognized in time so that timely decisions and counter actions can be taken. In other words the COP should contain sufficient actionable information to be retrieved on demand. To achieve situation awareness in a wide area, persistent surveillance, background intelligence, and multi-source data analysis are required.

3 Picture compilation and situation awareness

Picture compilation implies the collection of data about activities in the littoral followed by sense making of the data. Sense making can either be done by defining normal activities and searching for abnormal patterns or by recognizing signatures of normal and threatening processes. This results in the so-called Recognized Maritime Picture (RMP) which is defined as a composite picture of activities over a maritime area of interest. It contains tracks of vessels which have been evaluated with respect to the activity of the vessels. Analyses of the recognized processes in the RMP allow forecasting about future activities (i.e. situation awareness) including possible threatening activities. This information is input for the COP that contains actionable information for making decisions. In addition to the situational information from the RMP, the COP therefore contains all other information necessary for the decision making such as the position of own assets, capabilities and other relevant geospatial information such sea lanes, harbors etc.

3.1 Threat indicators, observables and situation awareness

We describe daily commercial or leisure activities in the observed world, such as fishing, trade and pleasure cruising, as process patterns. Processes can be recognized by sequences of situations which can be revealed by so-called indicators. Threats are processes that may occur and which are obviously not wanted. Situation awareness and threat awareness [1] implies here recognition of these processes before the unwanted situation has occurred. For example by producing a threat alert an operator and his supporting systems can become aware of an imminent unwanted situation. (i.e. the alarming situation, figure 1).

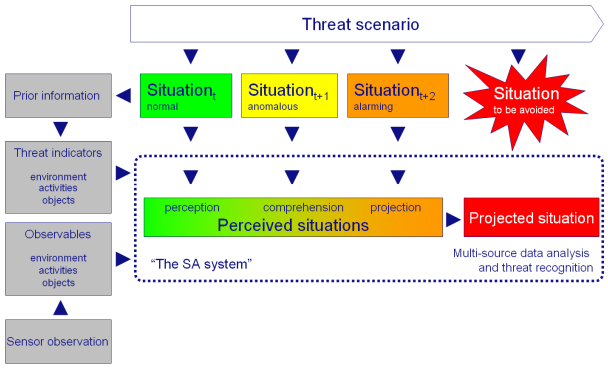

Figure 1. Diagram showing the relations between the scenarios and the sensor observations.

Examples of statements for threat indicators are: *a cargo vesselheading to a harbor other then the destination in the AIS message* in the case of smuggling or *a small vessel approaching with high speed* in case of a terrorist attack. In order to recognize the threat scenario, patterns of indicators, in which the threat is expressed, should be 'revealed' by the observables, which are elements of information suitable for describing the current situation and which can be directly derived from measurements performed by sensors. Sensor measurements yield data about the position, the characterization of the target, the track behavior, AIS information and the observation time.

4 Procedures for achieving situation awareness

Situation and threat awareness should be achieved by combining information describing the current situation expressed by *observables*, with signatures of threatening and normal processes expressed by *threat indicators* (see previous section). In our case the vessel is the object of interest in these processes and the requested information for the RMP is the *vessel mission*: trade or fishery etc (normal process) or smuggling, piracy etc (threatening process).

Observables involve statements which describe the current situation for which persistent surveillance, i.e. continuous tracking and tracing of vessels with observations systems, is a necessity. Therefore a broad suite of platforms equipped with sensors like, radar, AIS receivers and EO/IR systems is needed. Platforms comprise ships (both military as well as commercial), UAVs, satellites and VTS (vessel tracking services) ground stations. Examples of observables are statements about the size of vessels and ships (*large, small*), speed (*slow, high*), and track behavior (*loitering, stopped, and continuously ahead*).

Intelligence provides us with information describing the current context with respect to the vessel mission. This context consists of knowledge and information about the geophysical and geopolitical world (long term), information about current practices and trade activities (middle term) and information about recent activities of groups and persons and events (short term). Since the request for information (vessel mission) for the RMP is focused on the vessel, an object present in the spatio-temporal world, this current context should also be expressed in the same spatio-temporal world. For this purpose we have chosen the concepts: time, position, harbor and the vessel itself, for which intelligence should provide us *a priori information* with respect the possible vessel missions.

Figure 2 depicts the combination process where observables, indicators and mission a priori information are to be combined to determine which vessel missions are applicable for the vessel under surveillance. In het heart of the combination process, fusion methods are to be applied.

In the next section we discuss a number of possible fusion methods.

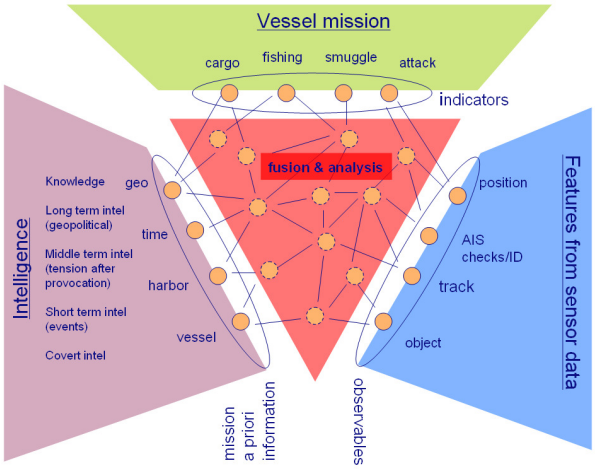

Figure 2. Information flows for sensor data and for intelligence, which have to be fused and correlated with the signatures of the vessel missions.

To build the picture and achieve situation awareness communication is essential through wireless broadband networks allowing for information exchange between assets and platforms. Data should be stored in databases in standardized formats so that members of coalition forces are able to retrieve and process the required information. An example is the Coalition Shared Database (CSD) developed by NC3A in the NATO/MAJIIC project (*www.nato.int/docu/update/2007/pdf/majic.pdf*).

5 Information fusion and analysis method

The fusion process requires that observables, indicators and mission a priori information and their interrelations are represented in a meaningful manner and readily accessible to the system. We need an information model that captures the context consisting of concepts and relationships that are relevant in our application domain, and that ensures consistency and a common vocabulary across the system components.

5.1 Using ontologies to capture the context

In order to be able to discriminate between normal and suspect and to compare various scenarios use of the context is inevitable. The context is the set of facts or circumstances that surround a situation or event. The context provides the background against which we can explain observations, and from which we can infer the most likely intent of a vessel. The context is given by procedural and factual knowledge about the world, about relevant objects and events, and their interrelations. The context can be expressed by a context model, an information model that contains relevant concepts and their relations, and encapsulates background and foreground knowledge required for understanding events. For example, in our maritime application domain, the context model should make it feasible to explain observations about vessels by VTS stations, and interpret messages from intelligence sources to explain their importance to the current situation. To this end, the context model needs to include all a priori knowledge necessary for the proper situation and threat assessment. We create such a context model via ontologies. We distinguish two types of ontologies: content ontologies, and situation ontologies [2], [3].

Content ontologies capture elements of interest in the application domain, such as known types of vessels and ports, and other domain-specific concepts. For most applications, one can use existing published ontologies to realise parts of the context model, such as ontologies on commonsense knowledge (WordNet, OpenCyc), geographic (GeoNames) and geopolitical knowledge, and many others. For the other part, one will need to resort to expert interviews and domain exploration to construct the required ontologies. Also, content ontologies serve to acquire data from information sources in a meaningful and consistent manner. Each element in a data source should be related to one or more concepts in the overall domain ontology, so that services can use information from sources in a consistent manner. In this way data sources containing information about a vessel but in different formats can be linked (e.g. databases using different ways to characterise the vessel type).

Situation ontologies capture a situation or series of states in the application space using concepts from the content ontologies. For instance, one might create an ontology that defines relevant geospatial vessel behaviours, such as the position relative to fishing grounds, or temporal patterns, such as arrival of a vessel at the scheduled time. We use such situation descriptions to characterise interesting behaviour patterns, and use those to establish our threat assessment process. The situation ontologies form the constituents for our search patterns of interest, e.g. the specific intents that we want the system to recognize. For example, a search pattern for *smuggle* might consist of a sensibly connected set of elemental patterns, such as *rendez-vous with another vessel mid-sea* or *the vessel type does fit the current location and time.* Colloquially, one might say that the situation ontologies define the observables and indicators to be used for intent recognition. See [3], [4], [5] for similar approaches. Figure 3 depicts the elements of the context model.

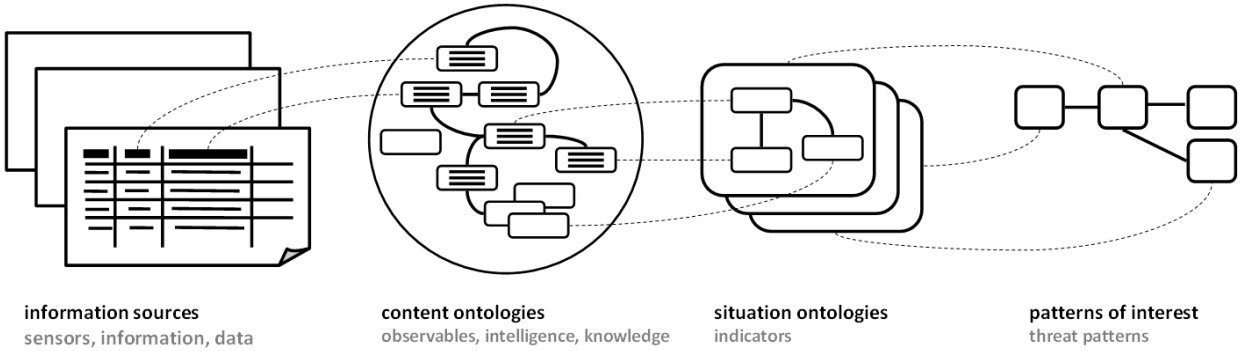

Figure 3. The elements of the context model. The content ontologies jointly form the domain ontology. Situation ontologies describe relevant situational concepts using concepts from the domain ontology. The situation ontologies form the basis for the definition of patterns of interest that may be used for threat assessment.

5.2 Exploiting the context model

Situation recognition requires that relevant relationships between threat scenarios, threat indicators, observables and mission a priori information are observed and assessed. The content ontologies provide semantic relationships; the situation ontologies define patterns that express how these semantic relationships lead to relevant indicators. Note that the initial ontologies are merely schemas, not the data itself. During run-time, the content ontologies are instantiated by available data. Ontology instantiation is a continuous process, which provides a basis for analysis of the current situation as well as for the building-up prior knowledge for future situation analyses. In figure 4 we present as an example a linkage diagram showing a subset of the semantic network covering both intelligence and sensor data. The right side of the diagram expresses the generic semantic relationships between concepts within the application domain (a priori defined knowledge). The left side of the diagram illustrates instantiated data (current information), in which observed vessel behaviour is related to concepts that are relevant for threat assessment.

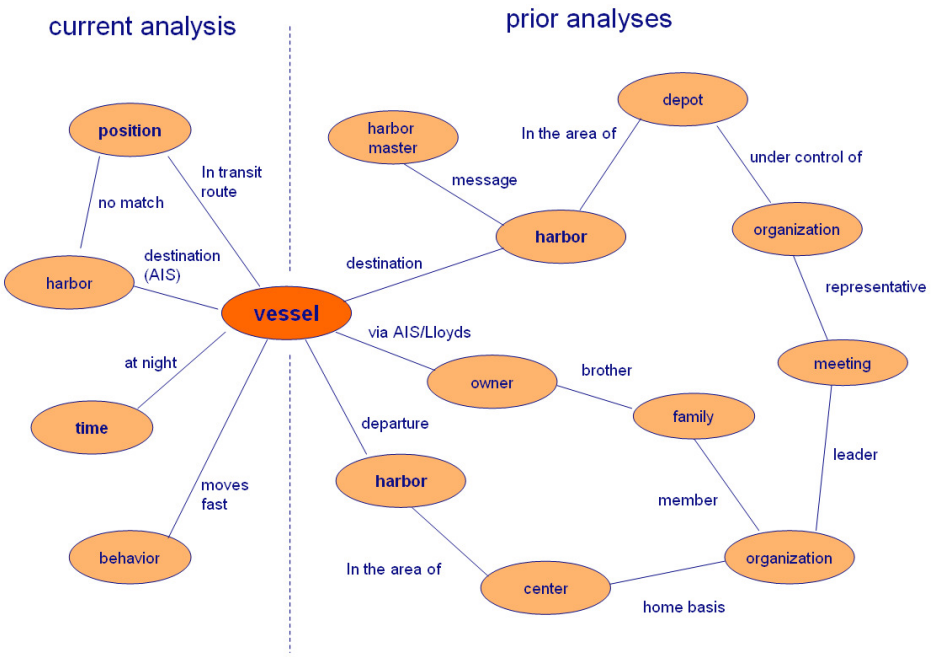

Figure 4. Linkage diagram derived from an underlying semantic network, with right the analysis in the intelligence domain and left the analysis in the sensor domain. Bold indicates concepts for which mission a priori information is available.

5.3 Assessment methods

The overall goal of the chain of processes is to provide alerts to vessel behaviours of interest. These behaviours can be expressed in patterns of interest, grounded in the context model. Basically, threat alerts can be produced by mapping threat patterns on the available, well-structured data. By exploiting the semantic relationships between data, we can conclude whether a threat alert should be issued.

Examples for assessment are probabilistic models such as Bayesian belief networks where prior information (based on previous observations or expert knowledge) is combined with the actual situation description (evidence) derived from the observations to produce the posterior probability of an hypothesis.

Typically a Bayesian belief network is based upon cause and effect relationships [6]. However, for application domains in which the causal nature of events is a regular subject of debate one may have to resort to an alternative approach. The Hypothesis Management Framework (HMF) was developed to enable decision support in such domains [7].

Effectively an HMF model is a Bayesian belief network that complies with a strict design pattern. Rather than pursuing a static model that represents the 'true' causality of the domain, an HMF model enables a flexible model that is easier to keep up to date with changes in the environment. In HMF competing hypotheses (in our case hypotheses about suspect versus normal vessel missions) are compared using sets of indicators. When an indicator is observed the posterior probability of each hypothesis is updated. The sets of hypotheses and indicators can be extended in a flexible way. Each indicator is independently related to each of the hypotheses. Newly added indicators or hypotheses will therefore not affect the integrity of existing prior knowledge in the model.

The HMF uses a so-called Naïve Bayesian Classifier topology. Such a topology has proven to be quite effective in getting good results [8]. Due to its structure it is also inherently robust for imprecise prior knowledge [9]: each indicator has a relatively modest influence on the posteriors of hypotheses. HMF is therefore suitable to assess relations between the indicators and observables defined by the situation ontologies and their instantiations, where the latter provide both the current situation description (evidence) as well as prior knowledge for the indicators in combination with the hypotheses.

In practice for (military) security operations there may not be enough data to instantiate the relations in the ontologies to produce evidence and prior information with sufficient certainty. In that case explicit results using the HMF and therefore actionable information cannot be obtained.

Therefore it would be worthwhile to get insight which indicators and which relations in the ontologies can be suitably instantiated. By selecting these indicators for use in the HMF more explicit results and more decisive information may be obtained. To get insight in the decisiveness of indicators, the nodes and links in the ontologies should receive a value of importance. This value is based on factors such as the prior knowledge, the information gain and entropy [10], [11].

The Bayesian approach described above is appropriate when the information obtained is uncertain which is often true in complex situations where hostile intent has to be inferred in the midst of overwhelming normal activities. When situations are less complex and threats are more explicit, the Bayesian techniques may be used in a more straightforward way to obtain actionable information. An example is given by the identification data combining process (IDCP) [12].

Another approach in less complex situations is the use of rules and decision trees for obtaining actionable information. In particular decision trees may be used, when indicators, which carry decisive information, can be determined. To get an insight in the decisiveness, information entropy or information gain can be used. When sorting the data based on the information gain, the more decisive indicators become visible. Decision trees provide a possible representation of one or more hypothesis. Decision trees can also play a role in a triage of the data to determine the importance of further investigation. Consider the case many vessels need to be assessed simultaneously. Based on a decision tree, decisions can be made which vessels need to be fully evaluated and which do not need further attention.

6 Case study

Above-mentioned procedures and techniques are to be implemented in a situation awareness support system. For this purpose we have developed a dedicated test bed. In figure 5 we show the processing chain and functional flow of the system.

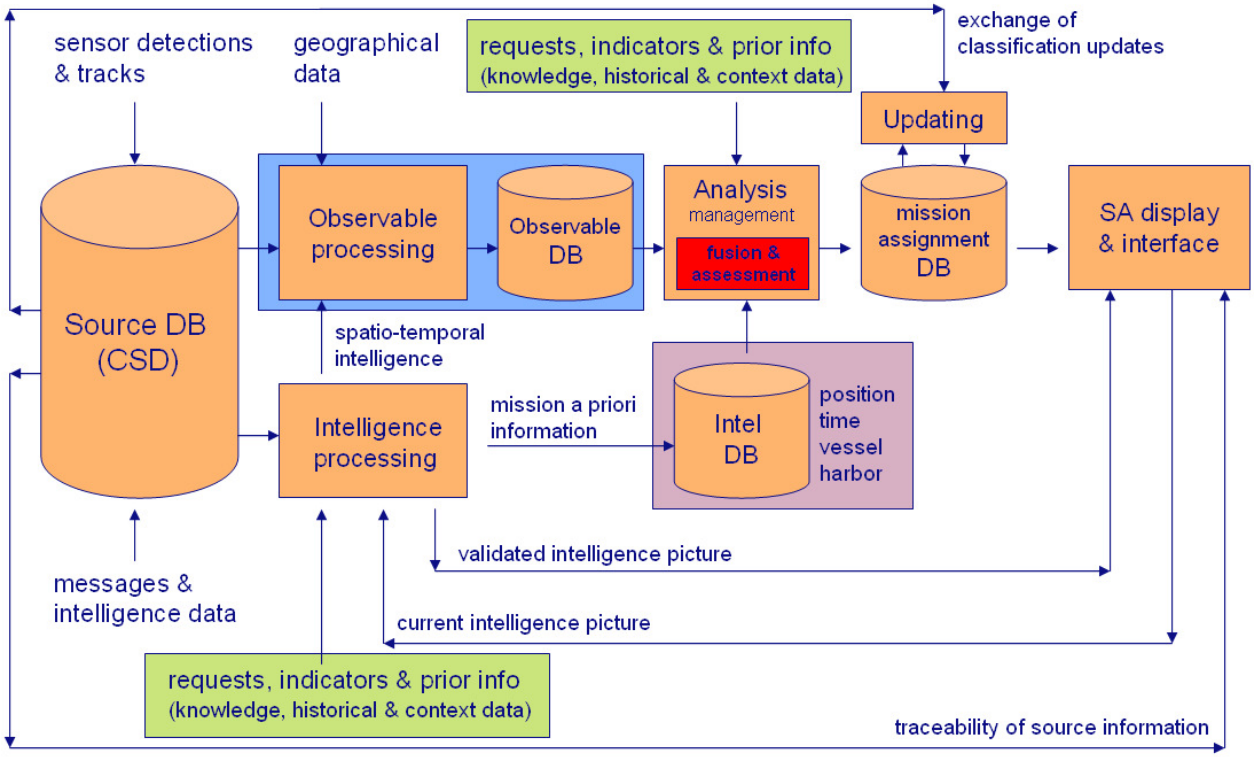

Figure 5. Processing chain and functional flow. Colours indicate the relation with the information flows in figure 2.

The system consists of several databases and modules :

1) *The source database* contains data collected by the sensors, intelligence reports and other collected information. In practice the coalition shared database (CSD) developed in Majiic can fulfill this role. Within the coalition the CSD comprises several databases in a distributed network where metadata about the contents is shared. On request the actual data can be retrieved.
2) *Observable processing* transforms the sensor data and tracks to statements about the detected vessel. It typically produces statements about the behaviour of the vessel such as *moves fast in a strait way*, *stays near coast*, *is in traffic lane,* etc. Inputs for the observable processing are geographical data about traffic lanes, coastal lines, positions of harbours, anchorage areas etc. Observables are produced and updated at regular time intervals and processing is mostly automated. These observables are becoming additional attributes to the tracks and are stored in the *observable database*.
3) *Intelligence processing* aims at producing (a priori) probabilities for the vessel missions with respect to harbour, time, position and the vessel itself. Inputs are the requests about the vessel missions to be revealed and long term a priori context information as well as knowledge from the physical world, historical data etc. Input data from the CSD, such as messages may not always be structured in standardized ways and intelligence processing is in practice not automated. The output information is updated when new information is becoming available and is stored in the *Intel database.* In this database the actual (a priori) probabilities about the vessel mission are stored for the position, time, vessel and harbours. The intelligence processing also produces a so-called validated intelligence picture which can be used as a layer for the RMP and COP. The information about time and position (spatio-temporal intelligence) is input for the observable processing
4) In the *analysis module* the actual fusion and analysis takes place. The red box indicates the *fusion and assessment module* where the techniques discussed in the previous sections are implemented. Input follows the three aspects depicted in figure 2: requests about the vessel mission, the sensor information flow specified by the observables and the intelligence flow specified by the mission a priori information. Output data are vessel mission assignments which are stored in the *mission assignment database*. In the *update module* the obtained mission assignments are compared with previous mission assignments to resolve conflicts and to obtain final results.

Once an assessment for the normal and suspect vessel missions (threats) are obtained the results can be used in the *situation awareness display module* and are combined with other layers such as the validated intelligence image, mapping and meteo data for the RMP. Also, threat alerts can be produced, which can be used for the COP and for decision support modules. Important item here is the traceability to the source database and information that caused the threat alerts, and access to other relevant information. The information is also made available for the intelligence processing module for future evaluations.

6.1 Scenario & simulation

To try out the system described above we have to supply the system with a continuous data stream. Since in practice data are incomplete we use simulated data. We use J-ROADS [13] for simulation of the entities, and sensor modeling of the VTS stations and sensors on board of military ships.

To build normal and suspect situations we use a maritime scenario in a sea strait with dense trafficking, and suspect behaviour can be introduced (see figure 5). We have adopted the following scenario for our case study. Two neighbouring countries are separated by a narrow strait, about 80 NM wide. In Troubledland a traditional government resides which is supported by a privileged minority of the population. In Badland a revolutionary government is advocating worldwide revolution. It supports terrorism in general and specifically the destabilizing groups in Troubledland. Because of tensions between the countries and conflicts in the past a UN peace keeping mission controls the traffic in the strait to prevent large scale smuggling of weapons from Badland to the destabilizing groups in Troubledland. There also exists the danger of terrorist attacks since Badland accommodates fanatic groups whose goal is to expel the international peace keeping force.

For the daily life scenario we adopted numerous vessels present in the strait, each with 'normal' missions. These missions are defined as follows: between the two countries local trade and smuggling of small items such food, clothes, and consumer goods etc is ongoing. Also regular transport of persons by ferries exists. In the middle of the strait a shipping lane is running where large oil tankers and cargo vessels (international trade) are passing. Fishing occurs adjacent to the coastlines mixed with pleasure yachts and cruise ships.

For each mission we have specified the type of vessel, behaviour, availability of AIS, and AIS content. In the scenario we also defined information about groups, individuals and their interactions, special locations and various events in order to be able to analyze intelligence. Information about the events is specified in messages, e.g. from harbors masters and other local agents. For monitoring purposes several VTS radar/AIS stations are positioned along the coast. The UN mission comprises three ships - two patrol vessels and one frigate - capable of monitoring the surrounding environment with their radars.

Using the system, scenario and simulation we can determine the applicability of the various techniques: ontologies, information certainty, HMF, decision trees & rules. At the time of writing this paper we are producing results which have to be evaluated before conclusions can be drawn.

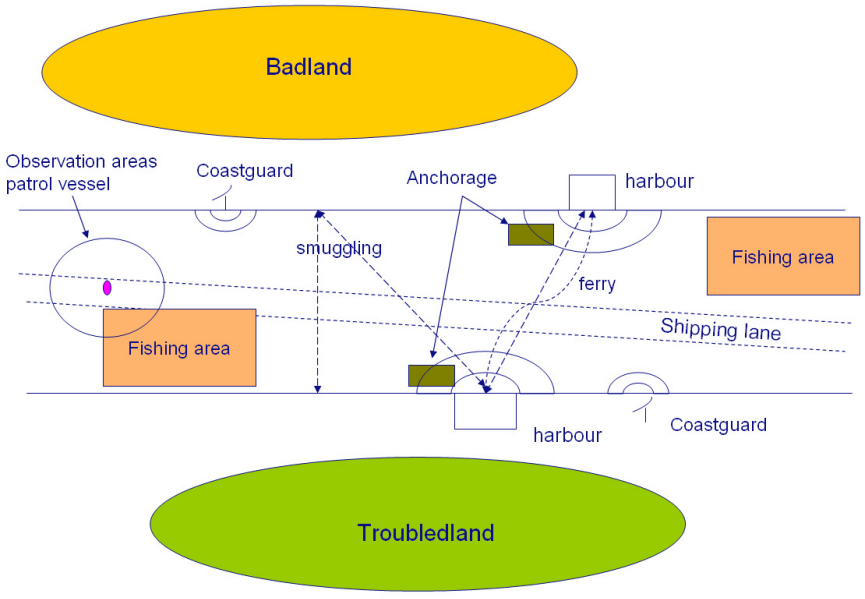

Figure 6. Schematic layout for the scenario

7 Discussion and conclusions

We have introduced and described a situation awareness support system focused on maritime security operations where sensor information is fused with intelligence data.

The fusion and analysis of the data for revealing suspect from normal behaviour is based on domain ontologies. A test bed allows the study of various exploitation and assessments techniques of the domain ontologies. Using an appropriate scenario and simulation of suspect and normal behaviour we are able to test the applicability of the various techniques.

In practice requests for on demand information and prior information will change and the system should be adaptable. Attention may shift to different suspect behaviours, or new insights into the intent of vessels may require the search patterns to be updated. This implies changing sets of patterns of interest and changes in the uncertainty propagation. This means that we need a method to dynamically update search patterns, and subsequently, update the exploitation and assessment methods to handle the new observables and indicators.

The introduction of additional observables and indicators may also imply new concepts in the ontologies. For different environments and different threats the system therefore has to be arranged or configured by human operators. A way forward is to identify those parts of the system which are suited for implementation of automatic adaption techniques. Guideline here will be operational practice.

References

[1] Endsley, M.R., (1995), *Toward a Theory of Situation Awareness in Dynamic Systems*, Human Factors 37(1), 32-64 (1995).

[2] Llinas, J., Bowman, C., Rogova, G., Steinberg, A., Waltz, E., & White, F. (2004). *Revisiting the JDL Data Fusion Model II*. Paper presented at the 7th International Conference on Information Fusion, Stockholm, Sweden.

[3] Kokar, M.M., Matheus, C. J., Baclawski, K. (2009). *Ontology-based situation awareness*. Information Fusion, volume 10, pp. 83-98.

[4] Baumgartner, N., Gottesheim, W., Mitsch, S., Retschitzegger, W., Schwinger, W. (2010). *BeAware! - Situation Awareness, the Ontology-Driven Way*. Data and Knowledge Engineering, Special Issue on Contribution of Ontologies in Designing Advanced Information Systems, volume 69 (11), Elsevier.

[5] Hage, W.R., Malaisé, V., Vries, G. de, Schreiber, G., Someren, M. (2009). *Combining ship trajectories and semantics with the simple event model (SEM)*. Proceedings of the 1st. ACM International Workshop on Events in Multimedia, Sheridan Publishers.

[6] Pearl, J., (2000), *Causality: Models, Reasoning, and Inference*, Cambridge University Press.

[7] Gosliga, S.P. van, Voorde, I. van de (2008), *Hypothesis Management Framework: a flexible design pattern for belief networks in decision support systems*". Proceedings of the 6th Bayesian Modelling Applications Workshop at UAI 2008, Helsinki, Finland, July 2008.

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

[9] Renooij, S., Gaag, L.C. van der, (2008), *Evidence and scenario sensitivities in naive Bayesian classifiers*, International Journal of Approximate Reasoning, Volume 49, Issue 2, pp. 398-416.

[10] Calmet, J. and Daemi, A., (2004) *From entropy to ontology*, in: Trappl R.: Proceedings of the Seventeenth European Meeting on Cybernetics and Systems Research (EMCSR 2004), 13-16 April 2004

[11] Cho, M., Choi C., Kim, W., Park, J., Kim P., (2007) *Comparing Ontologies using Entropy,* Proceedings of the 10th International Conference on Convergence Information Technology (ICCIT), 27-29 December 2007

[12] Kruger, M., Ziegler, J., (2008), *User-oriented Bayesian identification and its configuration*, Proceedings of the 11th International Conference on Information Fusion, Cologne, June 30- July 03, 2008

[13] Wiel, W., van der, (2006), *J-ROADS Air Defence Simulation Support during the 2006 JPOW IX Missile Defence Exercise*, NATO RTO, MSG-045-20