A Study of Anomaly Detection in Data from Urban Sensor Networks

Christoffer Brax[1] and Anders Dahlbom[2]

[1] Training Systems & Information Fusion,
Business Area Electronic Defence Systems,
Saab AB, Skövde, Sweden
christoffer.brax@saabgroup.com
[2] Informatics Research Centre, University of Skövde,
P.O. Box 408, SE-541 28 Skövde, Sweden
anders.dahlbom@his.se

Abstract. In many sensor systems used in urban environments, the amount of data produced can be vast. To aid operators of such systems, high-level information fusion can be used for automatically analyzing the surveillance information. In this paper an anomaly detection approach for finding areas with traffic patterns that deviate from what is considered normal is evaluated. The use of such approaches could help operators in identifying areas with an increased risk for ambushes or improvised explosive devices (IEDs).

Keywords: Anomaly detection, decision support, traffic flow analysis.

1 Introduction

In both peace keeping and peace enforcing missions, task forces mainly operate in asymmetric conflict environments where the situation most often can be described as being in a grey area between peace and war. Threats are usually camouflaged and hiding within the population and the regular activities of everyday life and the warfare is often carried out using terrorism, sabotage, IEDs, smuggling operations, etc. Task force security has become an increasingly important issue during missions in these environments, and new requirements are put on the technological support that is needed. It is not sufficient to detect the presence of an object in order to determine the threat that it might constitute.

A common tactic in asymmetric conflicts is the use of various forms of ambush attacks against the least defended elements followed by subsequent rapid movement away from the area of the attack [1]. In this way, attackers are only exposed for a very limited amount of time. Moreover, attackers carefully avoid open confrontation with larger and more heavily equipped forces.

In order to successfully identify these types of threats, suspicious object behaviors need to be detected and connected to imminent attacks or the preparation for them. This puts requirements on detailed information about detected objects along with robust processing of data received over time.

V. Torra et al. (Eds.): MDAI 2012, LNAI 7647, pp. 185–196, 2012.

Modern sensor systems are not designed to accomplish such surveillance tasks and it is left to human operators to detect and analyze these types of situations. Presented to an operator, all individual vehicle tracks in a city would however become unmanageable since it is difficult to maintain focus on more than a few tracks simultaneously. The operator, having limited cognitive ability [2], will also have trouble finding small changes in a situation. This might pose a problem in situations that extend over a long period of time with only small incremental changes. Additionally, there is a risk of operators experiencing information overload, which in the end can lead to poor threat detection. Only in exceptional cases can sufficient situation awareness in the ground theatre be achieved through conventional detect and track methods.

To address the situation awareness problem in urban environments, one has to either drastically limit the area of surveillance, or one has to look for objects that in some way deviate from a large background of similar objects (e.g. behavior, position). Anomaly detection is an interesting approach which pursuits the second alternative. The purpose of an anomaly detection function is to assist an operator by analyzing situation data to filter out important parts and give early warnings when suspicious behaviors are detected. When the function detects an anomaly it must be able to characterize it in such a way that the operator can easily understand the anomaly and make an informed assessment as to whether to monitor the object, take preventive action or to classify it as irrelevant.

Much of the focus in previous anomaly detection approaches in the surveillance domain has been on tracking and analyzing single objects to find objects that behave anomalous [3,4,5,6]. While this is important, it requires high quality tracking of the surveyed objects, something that is not always available in crowded urban environments.

This paper proposes a Gaussian anomaly detector which, in contrast to modeling the behavior of single objects, focuses on modeling the collective behavior of objects. This is carried out by constructing a model of normalcy based on the average flow and speed of objects in relation to geographical areas. Measuring the average flow and speed of objects does not require the use of advanced sensor systems with high quality tracking, and it is thus easier to carry out in urban areas. The proposed anomaly detector is evaluated using data from a simulation platform that simulates traffic in an urban area.

2 Anomaly Detection

One of the first fields to use anomaly detection was IT security where anomaly detection was used to build self-learning intrusion detection systems capable of detecting previously unknown viruses, trojans and break-in attacks [7]. Today it is also used in military and civilian surveillance systems. The concept of anomaly detection is, however, somewhat vague and there is no clear definition of anomaly detection is or even what constitutes an anomaly. Some argue that an anomaly is something that is known beforehand (i.e. can be described by a domain expert) but which seldom occurs [1]. Others argue that an anomaly is something

unknown that has not been seen before [6]. Anomalous objects are also referred to variously as outliers, novelties or deviations [8]. In this paper, we adopt the definition of anomaly detection and anomalies suggested by Tan et al.:

> "Anomaly detection is the task of identifying observations whose characteristics are significantly different from the rest of the data. Such observations are known as anomalies or outliers." [9]

A large variety of anomaly detection techniques have been suggested. Many of these have been specifically developed and tailored for specific problems and application domains, while others are more generic. Furthermore, anomaly detection has been the topic of a number of excellent surveys and review articles, see e.g. [10,11,8,12,7]. These articles mainly address methods which model data based on their statistical properties and use this information to investigate if the new incoming data originates from the same distribution or not.

During the past decade many different approaches for anomaly detection have been investigated in the surveillance domain. In [3] self-organizing maps are used together with Gaussian mixture models for detecting anomalous vessel traffic. This type of approach is also used in combination with interactive visualization in [13]. In [6] the focus is also on detecting anomalous vessel activity, however, they employ semantic networks composed of connected spiking neurons that are laid out in a grid over an area of interest. This allows for some degree of temporal information to be modeled. Other work focusing on detecting anomalous vessel behavior include the use of Bayesian networks [14,15], kernel density estimation and conformal prediction [5,16] and trajectory clustering [17]. Besides the maritime domain, work has also been carried out on anomaly detection based on e.g. video data. In [18], trajectory clustering is used for detecting anomalous traffic behavior and in [4] abstract state space modeling combined with Gaussian models is used for detecting anomalous behaviors in public areas.

Anomaly detection approaches in the surveillance domain have mostly focused on tracking and analyzing single objects to find objects that behave anomalous. While this is important, it requires high quality tracking of the surveyed objects, something that is not always available in crowded urban environments.

2.1 Anomaly Detection in Urban Environments

Urban environments are characterized by numerous closely spaced targets moving in rather confined spaces. In such crowded scenarios, the origin of observations is often highly ambiguous, meaning that it is a complex task to associate observations to new or existing tracks. In some situations, it is even impossible to determine the origin of the observation, no matter how advanced the tracking algorithm is. For this type of operational environment, additional functionality must be incorporated into the technical systems in order to analyze situation data and support the human operator.

It is thus interesting to look at techniques that do not rely on accurate tracks and advanced information about individual objects, but which instead make use of the collective behavior of objects as expressed by uncorrelated observations.

3 Gaussian Anomaly Detector

This paper investigates the use of a simple Gaussian normalcy modeling scheme modeling the collective behavior of objects in an urban setting using average speed and flow of objects. An assumption is that the collective behavior in an area where a threat is located is inclined to change. Moreover, road blocks and similar changes to trafficabillity will also have an effect on movement patterns. The idea is that threats possibly can be detected by finding changes in the traffic around them. It might however not be sufficient to construct a model over the complete area of interest, but rather to construct models for subspaces laid out in e.g. a grid, similar to [6]. In urban settings there also exists contextual information that can be exploited, e.g. maps of road networks.

The detection performance might however vary depending on the type of measurements that are used and on the type and degree of subdivision that is used. These are important factors to evaluate for deciding what kind of sensors that are appropriate and if contextual data such as road networks can be used for improving performance. Three important questions have thus been identified: (1) should we divide the area of interest using a simple grid or using additional contextual road segment information, (2) should we measure the average flow or the average speed of objects, and (3) how does the degree of subdivision affect the detection performance?

An anomaly detection system consisting of five components has been constructed for addressing these questions (figure 1 shows a schematic structure).

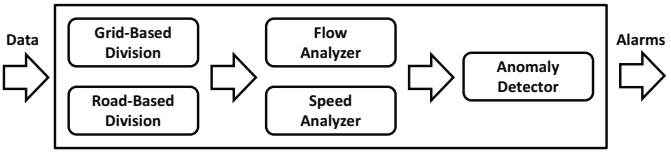

Fig. 1. Components in the anomaly detection system

The first two components handle the division of objects into subsets based on a grid or a road network. The next two components analyze the flow and speed of objects in each subset. The last component is the anomaly detector. The components at each step can be connected to any of the components at the next step, in order to easily evaluate various configurations of the system.

3.1 Geographical Division

The *Grid-Based Division* component creates a grid over the area of interest. The grid is defined by the upper right and lower left coordinates and the number of rows and columns. Each grid cell is represented by a geographical zone. When a new object is received by the system it will be assigned to the geographical zone

that contains the coordinates of the object. The use of a rectangular grid may result in e.g. roads being split into multiple grid cells and cells with few data points due to low coherence with actual geography.

The *Road-Based Division* component uses GIS data to create a number of geographical zones based on the segments in the road network. The GIS data defines a number of waypoints and information about how these are connected. The component assumes a 10 meter wide road between every pair of connected waypoints. When a new object is received by the system, it is assigned to one of the zones containing the position of the object. Note that there might be overlapping zones where road segments meet. This should not significantly impact the results of the experiments since the same zone is always chosen when there are multiple overlapping zones.

3.2 Flow and Speed Analayzers

The *Flow Analyzer* component measures the flow in an area by counting the average number of objects in the area over time. The analyzer handles multiple areas at the same time and these can be supplied by any of the two subdivision abilities. The flow analyzer could in a live system be implemented using some form of tracking or tripwire sensor.

The *Speed Analyzer* component measures the average speed of objects present in an area. Similarly to the flow analyzer, the areas can be supplied by either of the division abilities. The average speed is measured over all the objects in the area. In a live system, the speed could easily be measured using a doppler radar and it would not require any tracking.

3.3 Anomaly Detector

The Anomaly Detector component is responsible for the actual detection of anomalies. It can run in two modes: learning and detection. In learning mode, the component receives flow or speed statistics based on grid or road division. These statistics are saved and the mean and standard deviations are constantly updated for each geographic zone (grid cell or road segment).

In detection mode, a previously learned normal model is loaded and used to classify new data as normal or anomalous. This is carried out by inspecting if the present mean value is within n standard deviations (σ) of the mean in the previously learned model. If the new mean lay beyond this range, an alarm is sent to the presentation system. In this context an alarm consists of time, reference to geographical area and the actual deviation. In some cases, the standard deviation for a geographic zone is very small and therefore a parameter called minimum deviation (σ_{min}) is also defined. This threshold value defines a minimum required distance between the present mean and the model mean, for sending an alarm.

The last feature of the anomaly detector component is a sliding window that allows the detector to set which of the new pieces of data the detector should use when calculating the mean. The available options are: all data, last minute, last five minutes and last fifteen minutes. Setting the sliding window value to a

low value, e.g. last minute, will increase the reactivity of the detector but it will also make it more sensitive to variations in the input data.

4 Experimental System

In order to investigate the questions put forth in section 3 an experimental system has been constructed. This system has been built by integrating a number of existing software components and by extending their functionality to support evaluation of the proposed anomaly detection system described in section 3. Figure 2 shows the architecture of the experimental system.

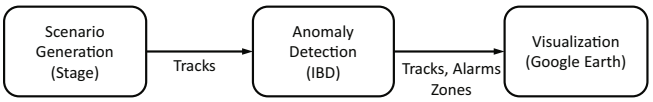

Fig. 2. High-level architecture of the experimental system

As can be seen, data sets are generated in real-time using Stage[1]. Individual tracks of simulated objects are fed into the Anomaly Detection System (implemented on the Intelligent Behavior Detector (IBD) platform [19]). The IBD routes the tracks to Google Earth[2] for visualization while at the same time using them internally for anomaly detection. The result of the anomaly detection are alarms that also are sent to Google Earth together with information about geographic zones that are used by the anomaly detector. To show information in Google Earth overlaid on the map, we use kml files[3]. The information in the kml file is regularly updated to reflect the current situation and it is generated from the current tracks, alarms and zones defined in the IBD.

4.1 Experimental Setup

A number of different experiments have been carried out to evaluate the usefulness of data-driven anomaly detection for detecting threats such as roadblocks and IEDs. Two simulated scenarios have been used for evaluation, where the first consist of normal traffic from an area of interest. The output from this scenario represents the training dataset that is used to train the anomaly detection algorithms. The second scenario is similar to the first scenario, but with a number of roadblocks/IEDs added to it. Due to the definition of roadblock, the generated traffic will automatically avoid these areas and instead use alternative routes to reach their corresponding destinations. The output from the second scenario has been used for evaluating different settings of the anomaly detection algorithm.

[1] More information about Stage can be found at `http://www.presagis.com/products_services/products/modeling-simulation/simulation/stage/`

[2] Information about Google Earth can be found at `http://www.google.com/earth/`

[3] Keyhole Markup Language, `http://code.google.com/intl/sv-SE/apis/kml/`

4.2 Simulated Road-Block Scenario

To create the two versions of the previously described scenario, it has been identified that Stage needs to be extended to fulfill the requirements in table 1.

Table 1. Requirements for creating simulated road-block scenarios

1. It should be possible to generate traffic that follows a road network.
2. It should be possible to dynamically spawn vehicles at multiple locations.
3. Vehicles should be able to take different routes though the road network to simulate different driving behaviors.
4. It should be possible to set the spawn interval for each spawn location.
5. It should be possible to set the parameters of spawned vehicles such as speed, type and initial heading.
6. It should be possible to turn the spawning on and off to simulate an uneven flow of new vehicles.
7. It should be possible to alter the road network and add roadblocks where traffic cannot pass.
8. It should be possible to extract the ground truth from the simulation as well as objects detected by simulated sensors.

The requirements are fulfilled by creating three new entities in Stage: (1) a spawner entity, (2) a spawn controller entity and (3) a vehicle entity. The spawner entity is responsible for spawning new vehicles and it can be configured to spawn vehicles at different intervals. The spawn controller entity is responsible for turning the spawner entity on and off at certain intervals, to simulate an uneven flow of traffic. The vehicle entity represents individual objects in the simulation. Each spawned vehicle is given a mission that it starts to implement as soon as it is spawned. The mission tells the entity how to behave. A number of different missions are defined such as *Go to beach*, *Exit the area east*, *Go to shopping centre* and *Go to beach and then to a specific parking lot.*

Figure 3 shows the area of interest and it also illustrates a single road block that has been placed outside a shopping mall in the evaluation scenario.

4.3 Experimental Process

The process that has been used for evaluation includes 7 steps:

1. Start the simulation with the normal flow scenario. Wait for the system to reach a stable state (usually five minutes).
2. Start the Anomaly Detector in learning mode and let it run for one hour.
3. Turn off the Anomaly Detector and save the normal model.
4. Turn off the simulation.
5. Start the simulation with the scenario including a number of flow interrupting elements. Wait for the simulation to reach a stable state.
6. Start the Anomaly Detector in detection mode. Wait five minutes to let the statistics settle and start measuring which anomalies are found.
7. Stop the anomaly detector and the simulation. Evaluate the results.

Fig. 3. The area-of-interest used in the simulation. The figure shows spawners and spawn controllers as blue triangles and the location of the road-block as a red circle.

A stable state means that there is a constant flow of traffic on the roads. When the simulation starts there is no traffic on the roads and vehicles begin to spawn at the spawning locations. After a while, when enough vehicles have been removed after reaching their destinations with new vehicles simultaneously spawning, the simulation will behave in a stable manner, i.e. the roughly the same number of vehicles are present in the simulation at any time.

5 Results

The first set of experiments is based on the ground truth from the simulation, i.e. the correct position for all objects at all times. This is not possible in real-world scenarios where real sensors must be used. Therefore, a simple sensor was implemented in the Stage tool for the second set of experiments. The sensor corresponds to the Saab SIRS 1600 short range radar sensor [20] that has a detection range of 1600 meters and a field-of-view of about 15 degrees. The SIRS 1600 can detect and track objects such as humans, cars, buses and trucks.

5.1 Results Using Ground-Truth Data

The default parameters in the experiments were $n = 2$ and $\sigma_{min} = 0$. In some experiments, the parameters have been altered in order to find any anomalies.

A total of eight experiments have been carried out in order to answer the questions in section 3. Each experiment has been carried out using the process described in section 4.3. Table 2 shows the results from the experiments.

Table 2. Results from experiments with ground truth data

Experiment	Division	Analyzed parameter	Result
1	Road	Speed	Works very well, no false alarms in normal data or during evaluation. An example of the alarms can be found in figure 4.
2	Road	Flow	Works very well, no false alarms in normal data or during evaluation.
3	Grid (10x10)	Speed	Hard to find the anomalies with the default thresholds. With $\sigma_{min} = 0.2$ some anomalies are found but the output are intermittent.
4	Grid (10x10)	Flow	No anomalies found. The road block is located on the border between two grid cells.
5	Grid (5x5)	Speed	No anomalies found. With $\sigma = 1$ and $\sigma_{min} = 0.1$ some anomalies are found as well as false alarms.
6	Grid (5x5)	Flow	No anomalies found. With $\sigma = 1$ and $\sigma_{min} = 0.1$ some anomalies are found as well as false alarms.
7	Grid (25x25)	Speed	Anomalies are found. Some intermittent false alarms.
8	Grid (25x25)	Flow	Anomalies are found. Some intermittent false alarms.

Should we divide the area of interest using a simple grid or using additional contextual road segment information?

Based on the results it is more efficient to use road segment information for subdivision, compared to using a simple grid. The optimal size of grid cells is however not obvious although the finest grid (25x25) yielded the best results in the experiments. The results from using the grid approach also resulted in unstable statistics in some cells. The normal variations were sometimes higher than two standard deviations which resulted in false alarms when feeding normal data into the anomaly detector. The road segment approach is therefore preferred if such data is available; otherwise the grid based approach can be used.

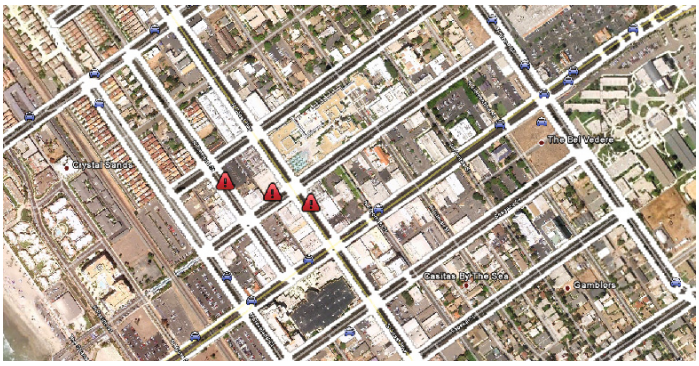

Fig. 4. Alarms from experiment 1

Should we measure the average flow or the average speed of objects?

In the experiments, the normal model based on both approaches performed similar. An advantage of using flow is however that the flows can be measured with simpler sensors.

How does the degree of subdivision affect the detection performance?

The experiments show that having too coarse a grid will decrease the ability to detect anomalies. The finest grid resulted in the most detected anomalies and the fewest false alarms.

5.2 Results Using Simple Sensors

Both the grid cell and the road based division of the geographic area were evaluated with the flow and speed analyzers. It was however decided that the grid cell approach should only be evaluated using the grid parameters that gave the best results, i.e. 25x25 grid cells. Each experiment has been carried out using the process described in section 4.3. The results are presented in table 3.

Table 3. Results from experiments using simple sensors

Experiment	Division	Analyzed parameter	Result
9	Road	Speed	Works very well, three anomalies detected in the vicinity of the road block.
10	Road	Flow	Works very well, two anomalies detected in the vicinity of the road block.
11	Grid (25x25)	Speed	Six anomalies are found. Four in the vicinity of the road block and one in each of the east and south entrances to the road network.
12	Grid (25x25)	Flow	Two anomalies are found. One near the road block and one at the west entrance to the road network.

The conclusion of the experiments with simple sensors is that the anomaly detection work almost as well as with ground truth data. This is however very dependent on the placement of the sensors and the number of sensors used. In the experiments, seven sensors were used to cover the most important roads in the area of interest. With fewer sensors, the performance would not be as good.

6 Conclusions

The detection of threats in urban asymmetric conflict environments, e.g. ambushes and IEDs, has become an increasingly important objective for increased task force security. Urban environments are characterized by numerous moving objects in crowded areas, making it difficult to only rely on detect and track methods. In this paper a Gaussian Anomaly Detector has been suggested for generating early warnings that can be used to assist human operators in the

detection of threats in such environments. The proposed anomaly detector focuses on modeling the collective behavior of objects of interest through the use of average speed and flow in relation to small geographic areas.

The initial experiments that have been carried out using a simulated urban threat scenario have investigated (1) different ways of dividing the area of interest, i.e. square grid cells or based on road-segments, and (2) if the average speed or flow was the best measure for modeling the normal behavior of a set of vehicles. The evaluation shows that, using the proposed anomaly detector, there is only small difference in performance between measuring speed and flow. It also shows that the use of contextual map information to divide the area based on road segments, yields more stable performance than using a grid cell approach. A conclusion is that road network information should be used if it is available; otherwise, acceptable performance can be achieved using grid cells.

An advantage of using an anomaly detector that operates on average speed or flow information is that it puts lower requirements on sensor systems and their tracking performance. It is not critical to have perfect tracking of all objects at all times; instead, it is sufficient to be able to measure the number of objects or the average speed of objects. This can be achieved using simple sensors.

Although it has been shown that the proposed anomaly detector can be used for detecting anomalies in a simulated urban scenario, more research and development is needed to achieve an operational system. The normalcy modeling scheme should be extended to handle more contextual information and to better capture variations in the data. Moreover, it is not enough to evaluate the feasibility using only simulated data. Data from a real sensor network deployed in an urban area should be collected and used for evaluation.

Acknowledgments. This work was supported by the European Defence Agency under the Defence R&T Joint Investment Programme on Force Protection; Contract No. A-0828-RT-GC Data Fusion in Urban Sensor Networks D-FUSE. The research has also been supported by the Infofusion Research Program (University of Skövde, Sweden) in partnership with Saab AB and the Swedish Knowledge Foundation under grant 2010/0320 (UMIF).

References

1. Ayling, S., Benchoam, D.: New antennas for new battlefields. In: Proceedings of the Military Communications and Information Systems Conference, Canberra, Australia, November 9-11 (2010)
2. Nilsson, M., van Laere, J., Ziemke, T., Edlund, J.: Extracting rules from expert operators to support situation awareness in maritime surveillance. In: Proceedings of the 11th International Conference on Information Fusion, Cologne, Germany, June 30 - July 3 (2008)
3. Kraiman, J.B., Arouh, S.L., Webb, M.L.: Automated anomaly detection processor. SPIE, vol. 4716, pp. 128–137 (2002)
4. Brax, C., Niklasson, L., Smedberg, M.: Finding behavioural anomalies in public areas using video surveillance data. In: Proceedings of the 11th International Conference on Information Fusion (Fusion 2008), Cologne, Germany, June 30 - July 3, pp. 1655–1662 (2008)

5. Laxhammar, R., Falkman, G., Sviestins, E.: Anomaly detection in sea traffic - a comparison of the gaussian mixture model and the kernel density estimator. In: Information Fusion, pp. 756–763 (2009)
6. Bomberger, N., Rhodes, B., Seibert, M., Waxman, A.: Associative learning of vessel motion patterns for maritime situation awareness. In: 2006 9th International Conference on Information Fusion, pp. 1–8 (2006)
7. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: A survey. ACM Comput. Surv. 41(3) (2009)
8. Patcha, A., Park, J.: An overview of anomaly detection techniques: Existing solutions and latest technological trends. Computer Networks 51(12), 3448–3470 (2007)
9. Tan, P.N., Steinbach, M., Vipin, K.: Introduction to Data Mining. Pearson Education, Inc., Boston (2006)
10. Portnoy, L., Eskin, E., Stolfo, S.: Intrusion detection with unlabeled data using clustering. In: Proceedings of ACM CSS Workshop on Data Mining Applied to Security (DMSA 2001), pp. 5–8 (2001)
11. Markou, M., Singh, S.: Novelty detection: a review-part 1: statistical approaches. Signal Processing 83(12), 2481–2497 (2003)
12. García-Teodoro, P., Díaz-Verdejo, J., Maciá-Fernández, G., Vázquez, E.: Anomaly-based network intrusion detection: Techniques, systems and challenges. Computers and Security 28(1-2), 18–28 (2009)
13. Riveiro, M., Falkman, G., Ziemke, T.: Improving maritime anomaly detection and situation awareness through interactive visualization. In: Proceedings of the 11th International Conference on Information Fusion (Fusion 2008), Cologne, Germany, June 30 - July 3, pp. 47–54 (2008)
14. Johansson, F., Falkman, G.: Detection of vessel anomalies - a bayesian network approach. In: International Conference on Intelligent Sensors, Sensor Networks and Information Processing, pp. 395–400 (2007)
15. Mascaro, S., Korb, K.B., Nicholson, A.E.: Learning abnormal vessel behaviour from ais data with bayesian networks at two time scales. Technical report, TR 2010 / Bayesian Intelligence (2010)
16. Laxhammar, R., Falkman, G.: Conformal prediction for distribution-independent anomaly detection in streaming vessel data. In: Proceedings of the First International Workshop on Novel Data Stream Pattern Mining Techniques (StreamKDD 2010), Washington D.C., USA, July 25, pp. 47–55 (2010)
17. Dahlbom, A., Niklasson, L.: Trajectory clustering for coastal surveillance. In: Proceedings of the 10th International Conference on Information Fusion (Fusion 2007), Québec, Canada, July 9-12 (2007)
18. Snidaro, L., Piciarelli, C., Foresti, G.L.: Fusion of trajectory clusters for situation assessment. In: Proceedings of the 9th International Conference on Information Fusion (Fusion 2006), Florence, Italy, July 10-13 (2006)
19. Saab AB: Intelligent behaviour detector. Internet, `http://www.saabgroup.com/Global/Documents%20and%20Images/Civil%20Security/Maritime%20Transportation%20and%20Port%20Security/IBD/ibd-ver1-090505a.pdf` (accessed January 28, 2012)
20. Saab AB: Sirs radar sensor. Internet, `http://www.saabgroup.com/Global/Documents%20and%20Images/Civil%20Security/Land%20Transport%20and%20Urban%20Security/SIRS%20200ITS/SIRS200ITS_downloads_Product%20sheet.pdf` (accessed December 28, 2011)