12th ICCRTS

"Adapting C2 to the 21st Century"

Automated Situation Assessment

Track 8

Dr John A H Miles
QinetiQ plc, UK
QinetiQ PTP, Portsmouth, UK, PO6 3RU
44 2392 312320
jamiles@qinetiq.com

Automated Situation Assessment

John A H Miles, Tony Edmonds, QinetiQ plc, UK

Abstract

Situation Assessment is a key element in Command and Control; it provides the cues (needs for action) and context (situation awareness) which enables effective decision making. The concept of automating such an apparently human process may seem alien but there is a case for it where human abilities are limited because of the sheer number of separate threads of reasoning, the length of time involved, and uncertainty in the information available.

Machine reasoning is able to track in real-time as many sequences of events as can occur, to remember everything required, to reason over periods of time, and to give consistent performance. What it lacks is the reasoning required for a particular circumstance and the ability to adapt its reasoning when circumstance change.

Having successfully applied knowledge-based techniques to generating real-time tactical pictures for warships and surveillance aircraft, QinetiQ has conducted research into automated situation assessment. The latest approach uses user-defined patterns and an initial exploration of the application of machine learning techniques.

The paper describes the techniques used and the results of experiments using both simulated and live data.

Automated Situation Assessment

John A H Miles, Tony Edmonds, QinetiQ plc, UK

The assimilation and collation of tactical data from local and remote sensors in a maritime Combat Management System (CMS) is a complex and difficult task, particularly in the littoral. Identifying a contact of interest from background traffic is becoming ever more challenging as new sensors and extra 'network enabled' sources increase the number of contacts reported to the CMS. A solution is, therefore, to consider how machine assistance could be developed and procured in an affordable manner to reduce the workload of the ops room team.

The command team has to make sense of a tactical picture which may comprise hundreds of system tracks representing contacts in the real-world. When assessing the picture, the operations room teams conduct two key activities:

- they observe the patterns of contact behaviour against prior expectations;
- they look for contacts that either do not conform to benign expectations or actively exhibit malignant characteristics.

The assessment process is in the first instance one of pattern matching to verify the benign or background traffic behaviour. The second case is to a degree the other side of the coin: when contacts fail to correspond to expected behaviour. When considering the pattern matching aspects of automated situation assessment (ASAS), the challenge is to show that machine assistance can exploit both prior knowledge about an operating environment (such as geography and local traffic conditions) and acquired knowledge from observed behaviour. This is to enable the command to assess the tactical picture more thoroughly, reliably and persistently than is currently possible. In so doing, the Commander should be able to make better use of scarce resources by focussing attention on unexpected contact behaviour, increasing the likelihood of mission success and reducing the number of errors stemming from incorrect identification and thus reducing fratricide.

An automated situation assessment capability offers the opportunity to free the user from the mundane, allowing him to concentrate on the remarkable; alerting the user to tactically significant changes in behaviour.

Combat Management Integration Support Environment (CMISE)

An existing test-bed known as CMISE was selected, with the support of the project sponsor, to host the pattern recognition rules. CMISE adopts a modular architecture, as shown within **Error! Reference source not found.** that already incorporated a module for automated tactical picture compilation, and elementary situation assessment and resource allocation modules. Each of these modules employs a generic blackboard Knowledge Based System

(KBS) methodology for representing and combining systematic rules to process and propose interpretations of tactical data.

The generation of the ASAS prototype as the enhancement of the situation assessment module (SAM) to include the pattern recognition rules was considered a low risk approach that enabled the new algorithms to be evaluated with the benefit of automated picture compilation to determine what reasonable performance can be expected in an operational system.

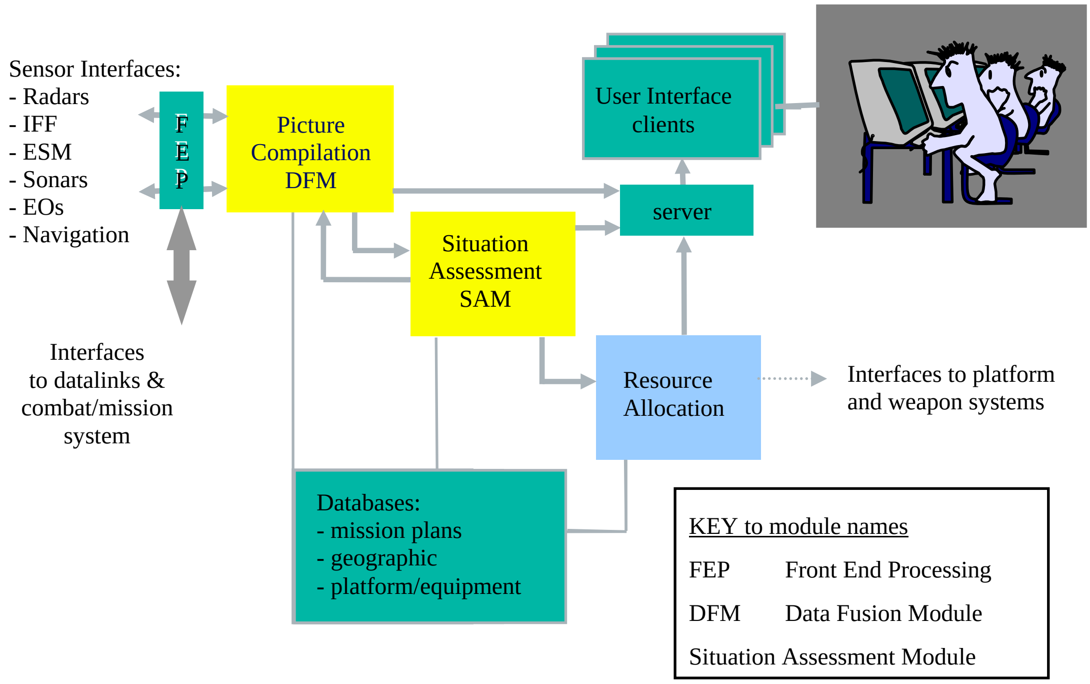

Figure 1 CMISE test-bed with knowledge-based components

<u>Requirements capture</u>

A series of knowledge acquisition meetings, with subject matter experts (SMEs) including technical specialists and experienced picture analysts, were conducted to identify anticipated operational patterns and consider a candidate concept of employment. Where operational patterns were identified, the attendees sought to identify suitable abstractions to represent these as systematic rules suited for machine processing.

<u>Test data</u>

The ASAS prototype was evaluated using tactical data recorded at a Shore Integration Facility (SIF). In addition, this 'real' data was supplemented by simulated data to represent operational conditions that could not be observed from the SIF. The simulated data, provided by the all environment real-time interoperability simulator (AERIS), represented a typical operational deployment.

Use of automated machine learning

In parallel to the ASAS prototype development, a study investigated the application of data mining techniques to pattern elicitation and specification to supplement the expert rule generation. These techniques are successfully employed within the commercial finance sector to provide fraud detection through pattern recognition within large volumes of data. This study sought to identify the merit of these techniques by evaluating their performance when processing the recorded data.

This study considered three specific data mining techniques and their potential to offer automated machine learning of tactical data patterns:

1. *cluster analysis* is used to identify groupings and regularities within a data set, allowing an analyst to visualise data and to explore the different patterns within it. This technique does not depend on pre-classification of the data;

2. *inductive logic programming (ILP)* is a type of supervised machine learning which, given a data set already pre-classified with 'labels', can be used to automatically derive a characterisation of the classification in the form of set of rules. Once derived, a classification rule can be used to classify any given instance;

3. *sequence detection analysis* seeks to identify and characterise repeated sequences of events within a data set.

This study has shown that data mining and machine learning tools and techniques can be used successfully to detect clusters and patterns in track data. Success is critically dependent on the quality and quantity of input data and on asking the right questions.

Evaluation of ASAS

The objectives of the evaluation phase were to:

- Measure the change in combat ID derived from pattern recognition;
- Demonstrate the generation of alerts on tactically significant events;
- Demonstrate the identification of spurious, benign and low priority information thereby allowing the operator to focus attention elsewhere;
- Consider the ability of machine assistance in deriving patterns;
- Evaluate the consequential picture quality changes;
- Evaluate the usability and usefulness of the system from an operator's perspective.

In recent years the single integrated air picture (SIAP) attributes have become increasingly recognised, recommended and applied to assessment of tactical picture quality. These attributes seek to measure the quality of information used to form a shared understanding of the tactical situation. The SIAP standards do not include a formal definition for combat ID amplifications, but nevertheless the SIAP approach was taken as a starting point for measuring ASAS derived picture quality changes.

Representing patterns

The model used in the ASAS prototype is designed to represent both 'simple' and 'interesting' patterns:

'simple' patterns: Neutral contacts are in general characterised by predictable behaviour:

- passenger airliners originate from airports, climb to cruising altitude and follow established routes;
- merchant ships usually follow well used routes in some cases delineated by shipping lanes, only altering course when constrained to do so by the International Rules for Prevention of Collision at Sea;
- airliners and merchant ships will all exhibit a 'departure' behaviour, related to a geographic point in space. They will 'transit' at a cruising speed along a route (which could be either tightly or loosely defined). They will 'arrive' when they move to another fixed point and stop;
- fishing vessels generally work in known areas at certain states of time or tide. They are often characterised by clusters of contacts moving in disparate directions, but remaining in a given area. A fishing pattern might, therefore, have a cluster of contacts remaining in company, with little overall directional movement.

'interesting' patterns: Some contacts are designated 'contacts of interest' since the command wants to know about them without any delay:

- some are the focus of an operation, typically when conducting indicators and warnings (I&W) patrols, or when conducting peace support or maritime interdiction operations (MIOP) tasks;
- some pose a real and imminent threat to 'own forces'. Such a contact could be a missile released from an aircraft, a closing weaving contact at Mach 2, or a torpedo in the water;

As a generalisation, the first category will be when a particular behaviour is seen in a defined geographical location. The second category represents threats.

The concept arrived at is that a pattern comprises one or more behaviour/area combinations, together with related identification and activity information. This is illustrated in Figure 2, where an aircraft is expected to keep to a particular route and then follow a racetrack. This could be an aircraft on combat air patrol (CAP) for example and can be represented as a single pattern containing two sequenced behaviours:

- a transit across area 1
- a loiter within area 2

If this is an aircraft on CAP, then the specific name of the aircraft may be known. On matching the pattern, the contact will be associated with that specific identity.

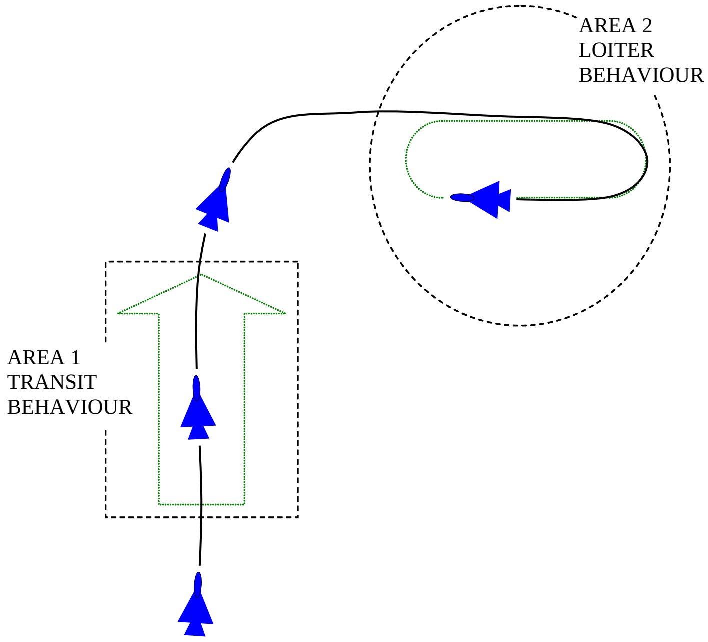

Figure 2 Example of building patterns from behaviour/area combinations

Thus the relationship between patterns, areas and behaviours, as illustrated in
Figure , is that:

- a pattern is composed of one or many behaviours
- each behaviour is associated with an area
- an area may be associated with multiple behaviours

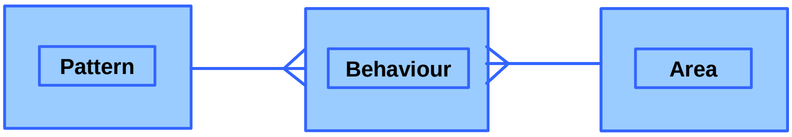

Figure 3, Entity Relationship between Patterns, Areas and Behaviours

Identity and activity information that can be associated with a pattern to support object identification comprises:

- platform identity (type, class or name)
- standard identity
- inherit group identity
- platform activity
- confidence
- whether the pattern is of particular interest

Behaviour and area combinations within a pattern may be sequential or non-sequential. Sequential behaviours are those that do not result in a pattern match if detected in the wrong order. Any existing pattern match is disassociated as soon as the sequence is broken.

Behaviour types

A prioritised list of behaviours was developed representing activity encountered during maritime operations. Some are benign, characterising routine maritime and air traffic flows, while other patterns are more clearly military and potentially malignant. Behaviours were prioritised according to their anticipated reliability, scenario independence, effectiveness and frequency of occurrence. The prioritisation determined which behaviours were to be implemented in the prototype.

A list of example behaviour types is as follows:

- transit;
- loiter;
- pop up;
- take off/leaving:
- speed/course/height change;
- landing/arriving.

Specific behaviour characteristics, such as speed or height bounds, are configurable by the operator and may be set up in advance, or changed while the system is running.

Areas

Every behaviour is associated with a tactical area. Tactical area types include:

- fixed and relative circles
- fixed and relative sectors
- corridors
- polygons

Additional tactical area parameters such as direction, height and time of validity support the ASAS architecture.

Pattern matching

The evidence required for matching a contact to a pattern depends on the behaviours that make up the pattern. Different rules were devised for each behaviour type. There is a balance to be struck between reducing the time taken to make a match and keeping the risk of error to a minimum.

When a contact first shows evidence of a behaviour, a *tentative* association is created. This is not considered to be sufficiently reliable to deduce platform identity or activity, but is the first indication of a match. When sufficient evidence has accumulated, a *confirmed* match is made, and the associated platform identity and activity are published to other modules. Tentative links between contacts and patterns may be displayed via the CMISE human computer interface (HCI) enabling the user to see to what patterns of behaviour the contact may be corresponding.

It may be possible for a ‘pattern break’ to occur when a pattern has matched to a contact and there is a subsequent change of behaviour. A pattern break will only occur if the behaviour is deemed not to have ‘completed’. Some behaviour types are considered to ‘complete’ at the point at which the confirmed match is made, an example of this is a loiter behaviour. Such behaviours can never be broken, this means that we want to retain the information that a contact has displayed loitering behaviour at a point in its history, regardless of subsequent behaviour. In other cases, such as the transit behaviour, a confirmed match can be made before the action has ‘completed’, e.g. before the contact has left the transit area via the exit boundary. If the contact deviates from the transit route and exits in the wrong direction the pattern match is broken. This is illustrated in Figure 1.

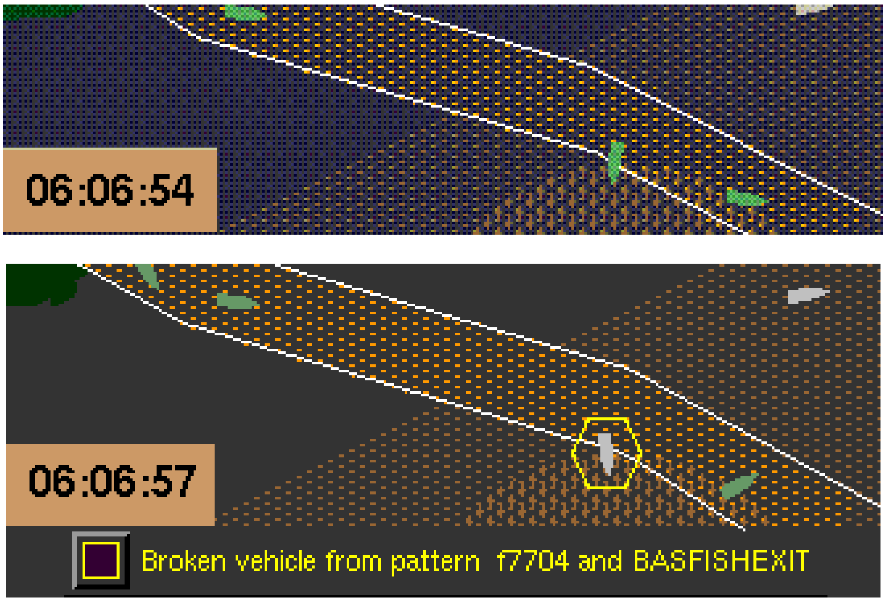

Figure 1 An example of a contact breaking from a transit behaviour

The operator is also able to manually associate/disassociate a contact with a pattern.

<u>HCI</u>

Facilities were provided to enable an operator to set up any number of patterns, behaviours and areas. New and changed patterns entered by the operator may be recorded, enabling them to be reused in similar situations.

Additional CMISE windows were developed to allow the user to list the following:

- all patterns that exist in the system;
- all patterns tentatively associated with a contact;
- all behaviours associated with a particular area;
- all patterns associated with a particular contact (confirmed links);
- all contacts associated with a particular pattern (both tentative and confirmed links).

Additional HCI functionality also allows the user to:

- click on the labelled plan display (LPD) to select points, areas and contacts when setting up patterns as an alternative to typing in latitudes, longitudes and identification numbers;
- highlight all contacts on the LPD that are associated with a particular pattern;
- highlight on the LPD the contact associated with a specific alert message;
- filter the LPD to show only those contacts that are associated with patterns that have been deemed ‘of interest’.

When defining a pattern an operator is able to request HCI alerts:

- if a contact is matched to a pattern
- if a contact breaks from a previously associated pattern

When an alert is raised the associated contact is highlighted on the tactical picture display. An example of an alert on breaking a transit behaviour is illustrated above in Figure 1.

<u>Contribution to ID assessment</u>

If a contact’s observed behaviour matches a pattern, then the identity and activity information associated with that pattern provide an additional source of evidence to the auto-identification algorithms within the CMISE data fusion module.

<u>Using ASAS</u>

The use of an ASAS application can be described under three life cycle stages:

- pre-deployment issues correlate to the preparation and use of ‘prior knowledge’;
- in theatre issues revolve around the exploitation of recently acquired knowledge and received intelligence briefs;
- post event analysis reflects how acquired knowledge and operational experience can be validated and fed back to enhance the store of prior knowledge.

The whole life cycle is captured in Figure .

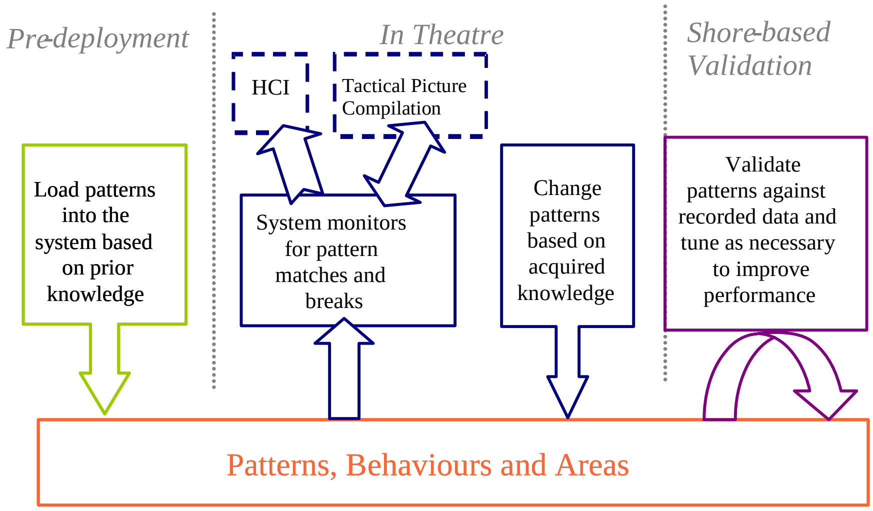

Figure 5, The life cycle of Patterns

<u>Human factors</u>

An ASAS capability should work on the user's behalf, continuously testing the latest system track data against a set of pattern definitions. Patterns can serve to identify routine movements of aircraft and shipping, allowing the user to focus his attention elsewhere. Patterns also have the potential to alert the user to unusual or threatening behaviour, or deviations from expected paths.

The user should be able to define a pattern as an area-behaviour-identity combination, thereby telling the system that if a contact is seen to conform to (or deviate from) a behaviour within a particular area it should be associated with a set of identity information. The intention here is to abstract the user's decision making rationale into a set of rules that may be followed by a machine. The logic behind any identity association made by the system should be clearly visible and intelligible to the operator.

Alerts are associated with patterns on a priority basis. For example, when a contact is matched to a low priority pattern (such as a ferry route), there may be no need to alert the user. Conversely, when a contact is matched to a high priority pattern (such as a manoeuvre presaging a missile release), the user needs to be alerted. When a contact breaks a pattern, an alert may be raised. Prioritising patterns is also the basis for 'clearing the undergrowth', allowing the user to concentrate on potentially malignant contacts. Various techniques, such as filtering or highlighting, may be employed for focussing attention away from low priority contacts and towards contacts of interest.

The clustering also enables operators to have their attention drawn to a contact with a significantly different behaviour characteristic. Three 35 kt contacts show up in the fishing fleet as suspected FIACs in Figure below.

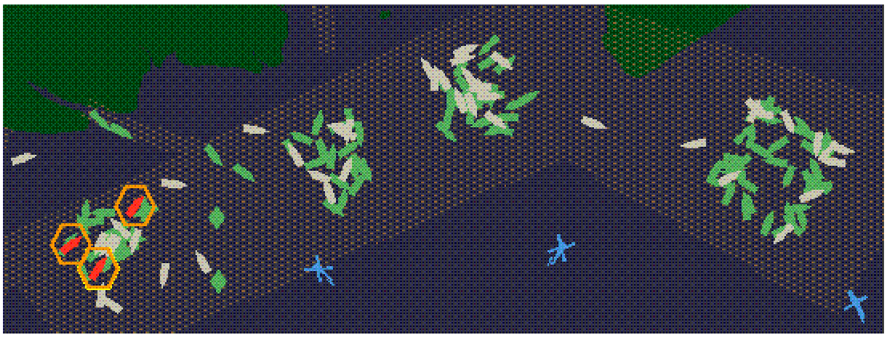

Figure 6. High speed contacts flagged within fishing vessel cluster

Live data trials

To determine how well the whether the ASAS functionality is likely to perform when driven with real data, performance was evaluated using recordings from a land based test facility with radars and other sensors.

Case 1: Airfield monitoring

Operationally monitoring airfields for aircraft taking off can be a laborious operator task. Given a significant number of such airfields in the battle-space, automated assistance for this task is likely to be of great value.

The concept is that ASAS monitors airfields, tagging all tracks of aircraft taking off with the airfield identity as their origin and generating alerts as required. The origin tag remains with the aircraft track, allowing the user to select all aircraft that have taken off from a given airfield, and provides a contribution to identity and/or activity.

Evaluation of the prototyped ASAS functions to carry out this monitoring task with real data requires a large set of aircraft take off examples. Busy commercial airports provide a significant number of movements in a practical 4-hour recording and replay period; London-Gatwick (LGW) was chosen as the test case as it is reasonable visible from the radars used and aircraft take off from there every few minutes.

With the visibility and the typical climb rate of commercial jets aircraft are likely to be tracked in under a minute after leaving the runway. An area was set up centred on the runway to capture tracks of aircraft taking off. A 'pop-up' was used with speed criteria.

To determine the setting for these speed criteria, track speeds within the vicinity of LGW were analysed – see Figure 2. The speed profile shows two distinct peaks; the one at around 110m/s (220kts) mainly corresponding to aircraft taking off and the other at around 220m/s (440kts) corresponding to aircraft flying through the area at higher level.

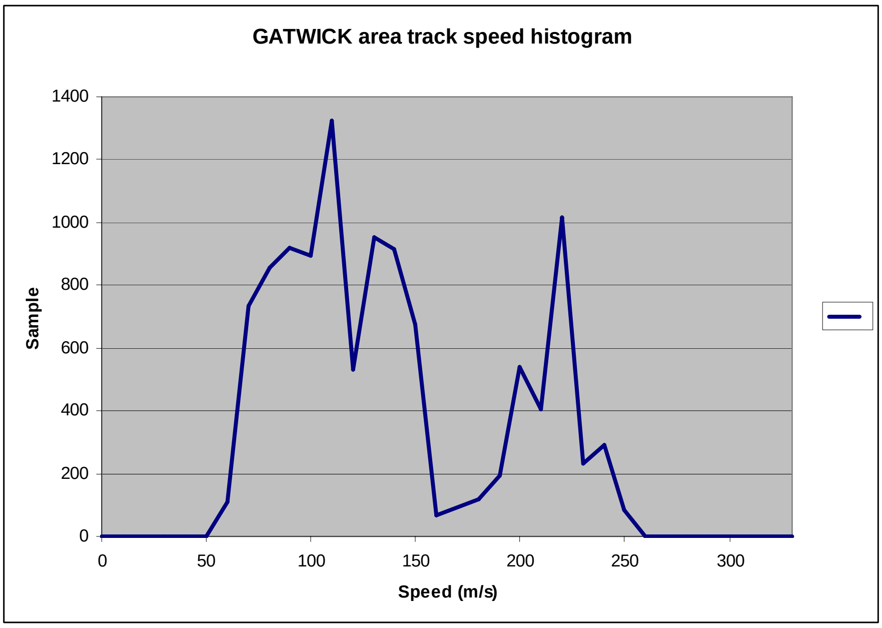

Figure 2 Speed profile for tracks in the vicinity of LGW

To evaluate the performance of ASAS with this set up, the tracks were analysed to identify every track representing an aircraft leaving LGW. Finding which tracks came from LGW was not difficult when examined retrospectively one at a time as the flight paths and speed profiles are significantly different to other air tracks starting in the vicinity.

In this sort of situation the reason to keep an area small and its attached behaviour as specific as possible is to avoid matching with other traffic which happens to be flying in the vicinity. However some aircraft are not being reported until later after take-off than estimated from line-of-sight considerations – an illustration of the variability that occurs in real radar detection and tracking.

To capture the majority of the take offs the area was set to a wide sector (±100°) with outer range at 12km as shown in Figure 3. This sector covers the tracks that do not start until after the aircraft have turned east or, in one case south, as well as those that start further west.

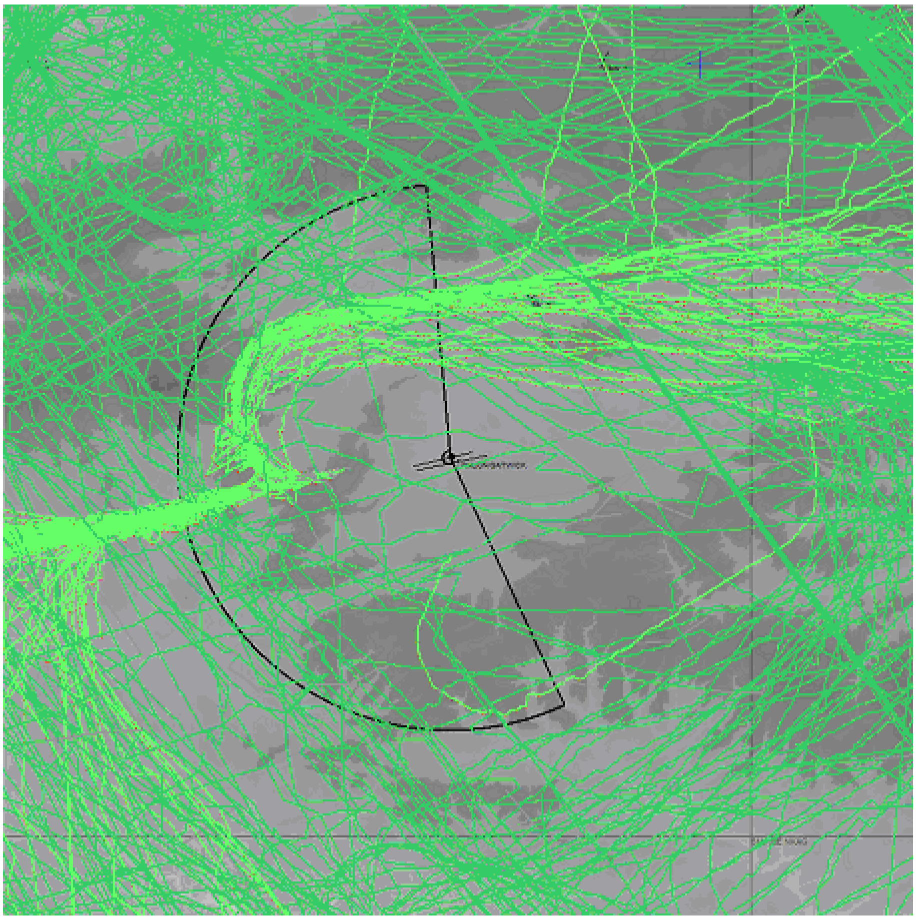

Figure 3 Area for capturing LGW departures

Using this set up ASAS was run on two days one week apart. 95% of take-offs were correctly identified and there were no false alarms.

These results demonstrate that robust performance can be achieved for airfield monitoring using the approach prototyped in ASAS and, importantly, with the available quality of air tracks from existing radars.

Case 2: Port/harbour monitoring

Although surface vessels proceed more slowly than aircraft and it is a less onerous task for an operator to monitor traffic emerging from a port or harbour, the vessels are usually more persistent in the area of interest which makes historical knowledge of behaviour and continual monitoring of behaviour more important. A system which picks up emerging vessels, constantly monitors their behaviour and retains the knowledge gained for easy access could be advantageous. Also it may be necessary to monitor a substantial length of coastline – not just recognised harbours because small fast craft require little in the way of shore facilities to operate; again this raises the potential value of automatic monitoring.

For ASAS evaluation purposes the concept is to monitor harbours and tag all vessels leaving with origin; then to continue to monitor traffic behaviour to identify benign vessel movements and generate alerts on vessels not conforming to expected routes and behaviours. Harbours in the Solent (UK channel between south coast mainland and Isle of Wight) provide many examples of vessel movements and AIS data provides the exact identity of most. Published time-tables also provide the times of departure of scheduled ferry services. It should be noted, however, that the number of movements at a Port in a 4-hour period is few compared with the equivalent period at a busy airport and hence performance measurement more difficult in the surface environment.

Portsmouth harbour (UK) was identified as a suitable local harbour for evaluation. Portsmouth has an interesting mix of traffic; it includes regular sailings of car ferries, fast catamaran ferries, and hovercraft to the Isle of Wight plus occasional appearances of pilot boats; representing about four movements per hour.

This situation can be likened to a harbour where there are benign commercial activities but also the possibility of fast attack/pirate craft emerging. Of the four types of vessel present in the recorded data three have predictable routes (the ferries) whilst the pilot boats are less predictable both in timing and the route they take to meet up with ships entering Portsmouth or, more often, Southampton. For ASAS evaluation it was, therefore, decided to treat the pilot boats as the vessels of interest – e.g. fast patrol boats. Of further interest, and adding to the realism, is the fact that the hovercraft, catamarans and pilots all travel at over 20kts making simple discrimination based on speed impossible.

To capture the benign and interesting traffic, ‘corridor’ areas were set up at the harbour mouth and along the ferry routes as shown in Figure (hatched areas). An area was also placed heading south along the deep water channel from Portsmouth to capture cross-channel ferries but there were no movements of these vessels in the recorded periods.

It took a few iterations to position the areas and select appropriate criteria to get the desired effect. A particular issue is that although speed is a discriminating factor for some vessels once they are well clear of the harbour, it is not at the point at which they need to be detected – near the harbour entrance where speed limits apply.

The figure shows a vessel emerging with suspect behaviour attached (red). Windows on the right hand side provide explanation of the supporting track data and contributions to identity including the pattern(s) matching – in this case the ‘pilots’ pattern. Coloured bars show the strength of the evidence determined from the type of pattern matching the contact.

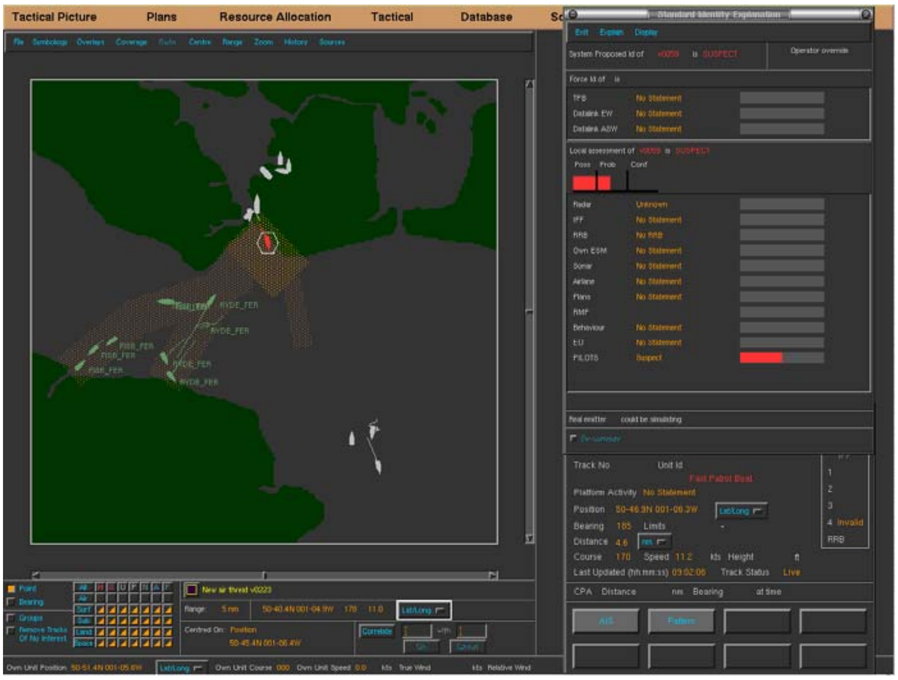

Figure 9 Vessel emerging from Portsmouth initially indicated 'suspect'

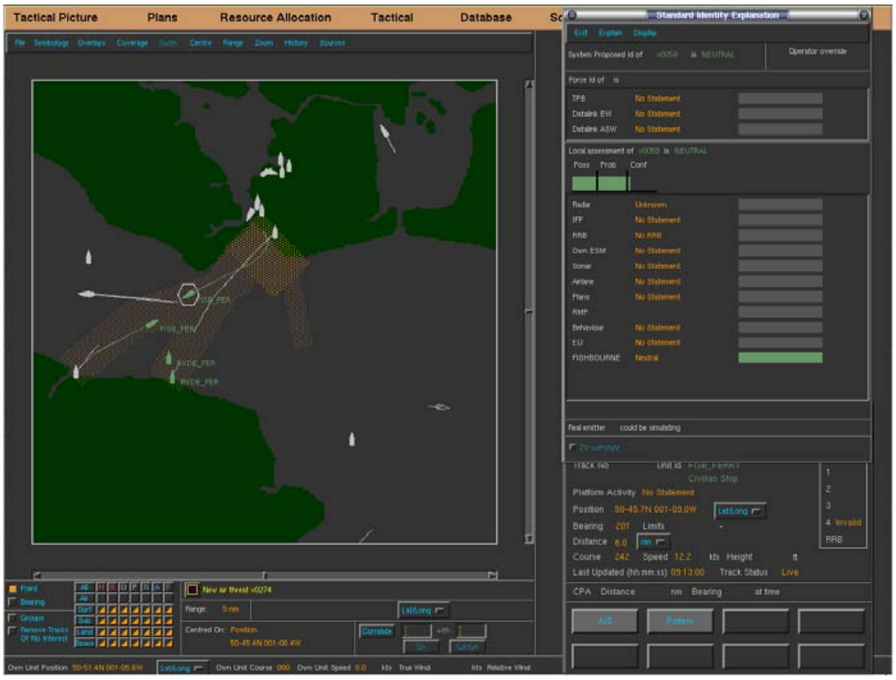

Figure 10 Vessel matching with just one ferry route

Later the vessel emerges from the overlapping areas at the harbour mouth and now has a unique match with the Fishbourne ferry route pattern; this in conjunction with speed criteria enables its route identity to be assigned. It will now retain this identity provided that it remains within the corridor for this ferry route and its speed stays within the limits set for this ferry.

The above sequence shows how a benign vessel is identified and monitored. Later in the run a pilot vessel emerges; this sequence of events is illustrated below in Figure 4 and 12.

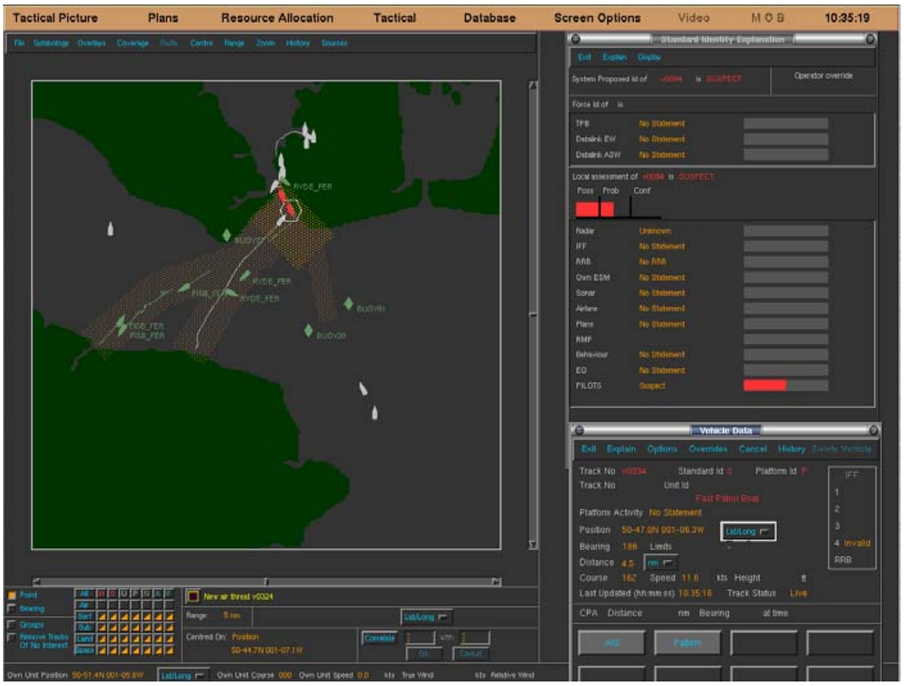

Figure 4 Two 'suspect' vessels emerge from Portsmouth harbour

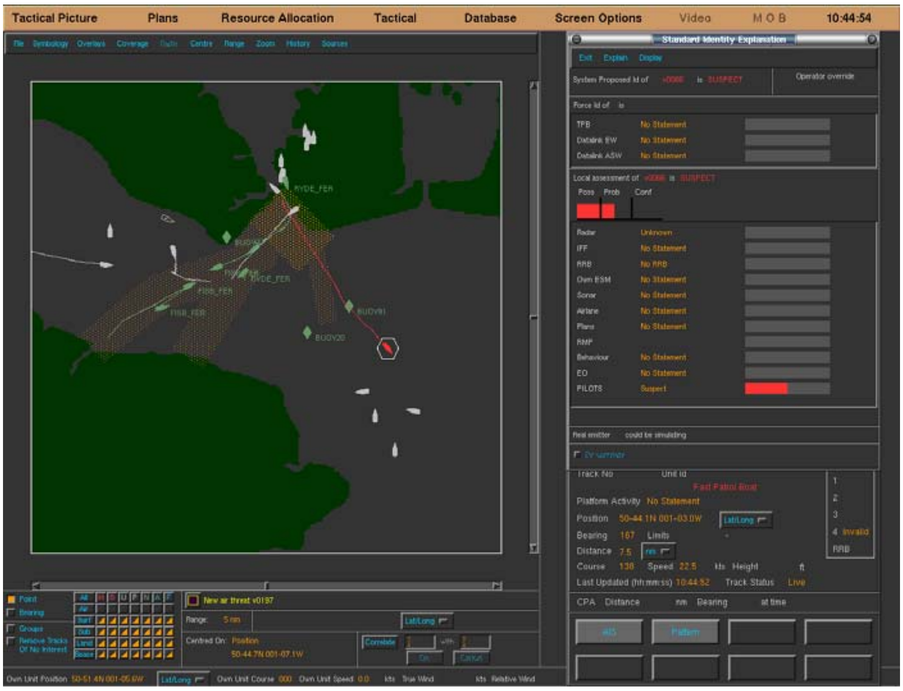

Figure 12 Vessel retains identity of 'suspect', 'fast patrol boat'

The above case illustrates plausible operation but has highlighted the possibility of refining the models of vessel behaviour particularly in terms of speed changes within the patterns so that better discrimination of vessels can be achieved. It also indicates that it would be useful to have a mechanism to generate identity or alerts if a vessel was not matching with any expected pattern – this would have allowed an easier and more effective approach to the Port monitoring case.

Conclusions

The ASAS investigation outlined a number of potential technical solutions and subsequently prototyped user programmable pattern recognition as a mechanism whereby a contact's behaviour is compared to its location and other intelligence to deduce what it might be. The ASAS prototype was exercised against plausible operational examples, some taken from synthetic data representing peace enforcement tasking, others taken from recorded live data from a Shore Integration Facility (SIF).

The capability has been shown to provide robust automatic monitoring and identity assessments given reliable track data and appropriately configured patterns, behaviours and areas. The evaluation from operators was positive, identifying several mechanisms whereby workload would be reduced. There were also plausible cases where behaviour of operational significance would be recognised more quickly thus reducing the risk of mission failure.

Benefits include

- Reduced operator workload;
- Alerts to focus attention on contacts of interest;
- Aid to more effective use of resources;
- Improved Situation Awareness;
- Applicable across land, surface and air environments.

To develop a full solution, additional work is required to identify priority operational cases with corresponding test data and to develop tools to assist the user in pattern configuration. Once achieved, the modular ASAS application would be available for integration into a CMS for seaborne user evaluation.

Data mining techniques could potentially be used to tune pattern criteria and may even offer the possibility of 'on-line' pattern learning in theatre; initial work has demonstrated the principle using a sample of techniques on some simple examples, but more work is required to explore other techniques and apply them to operational cases.