12th ICCRTS

"Adapting C2 to the 21st Century"

Automated Situation Assessment

Track 8

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

Automated Situation Assessment in a Maritime Combat System

Dr John A H Miles, Tony J Edmonds
QinetiQLtd, UK
+44(0)23 92312460
+44(0)23 92312386
jamiles@QinetiQ.com
ajedmonds@QinetiQ.com

Abstract

Situation Assessment is a key element in Command and Control; it provides the cues (needs for action) and context (situation awareness) which enable 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 circumstances 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 study included the implementation of user-defined patterns and an initial exploration of the application of machine learning. Techniques and experiments using both simulated and live data are described.

1. Introduction

QinetiQ has a long history of research into data fusion in maritime combat systems. JDL Level 1 data fusion, automating tactical picture compilation, has been evaluated at sea in a number of trials in recent years. The most recent tranche of data fusion research has been conducted in 2005/6 for Director Equipment Capability (DEC) Above Water Effects (AWE) in the UK MoD, to demonstrate the feasibility of level 2 fusion, automating situation assessment in a maritime combat system. This paper is a rapid walk through the thinking that shaped that programme of work and a presentation of key results and findings.

2. Background

Maritime Combat Management System (CMS) Process

A warship at sea typically has radar and associated sensors feeding data into a command team, whose job it is to make sense of the data, and adopt the appropriate course of action to accomplish its mission. The data processing task, illustrated in Figure 1 below, has been progressively automated as sensors have become more capable, and contact densities have increased. Raw radar is a thing of the past, and picture compilation is becoming increasingly automated. The poor relation that is the subject of this paper is Situation Assessment, where substantial benefit could be expected, were machine assistance shown to be practicable.

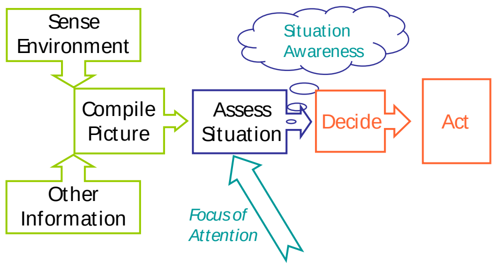

Figure 1. The generation of Situation Awareness within the Maritime CMS Process

Situation Assessment

It is quite usual for a surface escort to have in excess of 2½ thousand source tracks from which the picture has to be compiled. In coastal conditions, the command can be faced with hundreds of compiled system tracks from which sense has to be made. The command team is unable to give all contacts their full attention, so the question naturally arises: on just what should the operators concentrate?

The command's effectiveness is limited by the human's inability both to manage and interpret large volumes of data, and to concentrate on many concurrent events. The platform's effectiveness is similarly constrained by the need for users to provide assessments of the data for wider distribution around the network of military units.

The Hypothesis

Automated Situation ASsessment (ASAS) enhances C2 through rapid systematic reasoning and relentless exploitation of all accessible information; making results available across the network.

3. Current Research

Technical Approaches

Early attempts to capture situation assessment knowledge resulted in fixed rule sets for specific cases. More recently, study of operator working patterns, particularly around watch hand-over in the operations room focussed attention on two more powerful techniques.

During the watch handover, the users often describe broad swathes of activity either current or expected, before addressing more specific tracks or interesting features. The swathes set the context for wider appreciation and expectation. The behaviour of most contacts is characterised by generalisations rather than specific courses, ranges and speeds. Accordingly, the QinetiQ research has examined patterns of behaviour, be they either user-defined or more speculatively, derived from machine learning strategies.

User defined patterns largely mimic the human analysis. They are entirely under the control of operators and link together ideas of kinematic and other behaviours, geographical context and timing. Patterns can be derived from off-line data examined, in our case, in the laboratory after the event. They could equally be generated by post event analysis on a deployed platform and they offer the promise of generic libraries easing the process of pattern definition and knowledge transfer to platforms in theatre.

Data mining has been considered as a machine learning strategy. Various techniques could be employed, such as Cluster Analysis, Inductive Logic Programming, or Sequence Analysis. This more detailed examination of data offers an approach to identifying patterns of behaviour that are not readily observable by human operators but nevertheless could extract key features from operational data for subsequent exploitation.

What is a Pattern?

A pattern comprises one or more behaviour/area combinations, together with related identification and activity information. A behaviour is a period of kinematic activity; the movement of a contact over time. The time period is related to the nature of the behaviour.

Consider an aircraft, leaving an airfield. There is a departure activity, during which its behaviour is characterised by a particular rate of climb from low altitude, originating near an airfield towards a transit route. During departure, its heading is likely to vary considerably. Once it has reached its airlane, its behaviour will generally be 'straight and level' until it reaches a dogleg in the lane, or until it is ordered to a different height. The departure behaviour is not a function of aircraft course, but the transit definitely is. A series of behaviours can thus be proposed that relate to a number of segments of a typical vehicle journey.

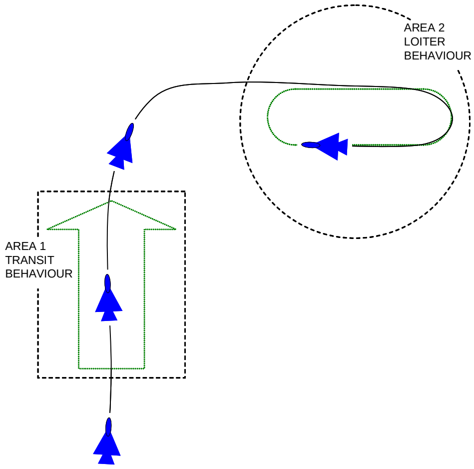

Figure 2. A combination of areas and behaviours may be used to identify the combat aircraft.

When a route and a loitering area are known, the fact than an air contact is following that defined pattern of behaviour can be a strong indicator to the identity of the aircraft, as in Figure 2 above.

Put simply, a Pattern can comprise several Behaviours that can be exhibited through several Areas. This is shown diagrammatically in Figure 3 below:

Figure 3. Relationship between Patterns, Areas, & behaviours

Pattern Implementation

Within our design and implementation, a number of patterns were postulated that had a range of attributes:

- Platform identity (type, class or name). An example would be pattern that recognises the Portsmouth to Fishbourne Car Ferry, or Flight VA021 from Gatwick Airport. It combines all that we know about the contact.

– Standard Identity. We can define patterns that capture friendly, neutral or indeed hostile activity. This is a valuable indicator when either ID is absent, or indeed in the face of spoofing.

– Platform Activity. Characteristic behaviour, such as fishing in a fishing area provides valuable corroboration. This could be used to generate a persistent attribute for the contact.

– Whether the pattern is of particular interest. There might be some patterns that, if detected, provide critical clues to an operational commander. Such a pattern match would therefore be of extreme interest. This could be used to 'watch' many airfields concurrently.

– Whether an alert is required on match/break. A major benefit from this machine assistance is that it is relentless and addresses all the data in the combat management system, irrespective of the command team's particular focus. It can therefore alert the operator to a configurable set of conditions whether they are matching or breaking patterns.

Postulated Behaviour

Each behaviour type that we have investigated has its own set of attributes. The types that we have implemented to date include:

– Loiter. A contact is remaining in a designated area with a minimal speed of advance in relation to the instantaneous platform speed. Various criteria can be defined for breaking out of a 'loiter'.

– Transit. A transit area is when there is a clear route, which might be curved or straight. The route has ingress and egress points where contacts can enter or leave the area. If a contact leaves the transit route through the 'sides', the pattern would be deemed to have been broken, as the contact was clearly not transiting in that case.

– Takeoff. Takeoff could be seen as a special case of 'leaving'. For the behaviour to be seen, there needs to be surveillance of the immediate runway area, or very near by. It is generally a climbing in conjunction with a manoeuvring characteristic.

– Approaching/leaving. An approaching or leaving behaviour would be linked to a specific geographical area or host platform. If it were an aircraft behaviour, it would also include climbing or descending.

– Pop up. This behaviour can be used for a remote airfield where the surveillance is at several hundred feet off the ground, or it could represent the launch of an air flight weapon.

Candidate Areas

Areas can be constructed with a number of attributes, specifically in relation to whether a matching contact is allowed to leave or join across specific parts. Area types include:

– Corridor. This area is closely allied to the transit behaviour.

– Polygon. Any shape and no specific aspects from which to leave or join the area.

– Circle. It can be fixed or moving.

– Sector. Usually truncated to reflect entry and exit through the inner and outer arcs.

Sources of Pattern Data

Geographic data can be found in the DAFIF database, giving all airports and airlanes, offering a ready source of material. Charts provide channels down which merchant shipping will travel. Separation zones are enforced by the coastguards. Ferry timetables provide intelligence on when regular traffic is expected down well defined routes. All this open source data can go a long way to clarify much of the background traffic in a typical littoral environment. For all the civil data there is corresponding classified information on what friendly forces are planned to do and intelligence on what hostile activity to expect.

There is also the possibility to parameterise those patterns identified from previous operations or platforms. It can then be passed on from 'corporate memory'.

4. Experimentation

ASAS functionality was implemented within the pre-existing Situation Assessment Module in QinetiQ's CMISE data fusion test-bed *Figure 4*.

The demonstration of ASAS feasibility was in two forms. There were live data trials using radar and associated self reported data at Portsdown Technology Park (PTP) in the south of England. There were also synthetic data trials using simulated military systems set in the Northern Gulf some 5-10 years from now in a fictitious setting. The live trials were to show that the algorithms could handle the sort of data quality expected from deployed sensors today. The synthetic trials were to illustrate military benefits in a realistic operational setting.

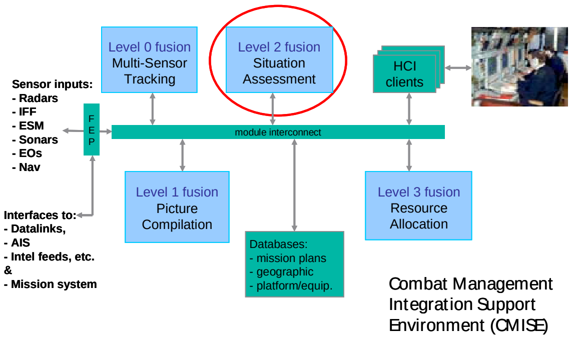

Figure 4. ASAS implementation within CMISE.

Live data trials

– 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 monitoring task with real data required 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) some 40 nmi from PTP was chosen as the test case as it is reasonably visible from the radars used and aircraft take off from there every few minutes. Given the distance, 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 5. The speed profile shows two general peaks; the one at around 110m/s (215kts) mainly corresponding to aircraft taking off and the other at around 220m/s (430kts) corresponding to aircraft flying through the area at higher level.

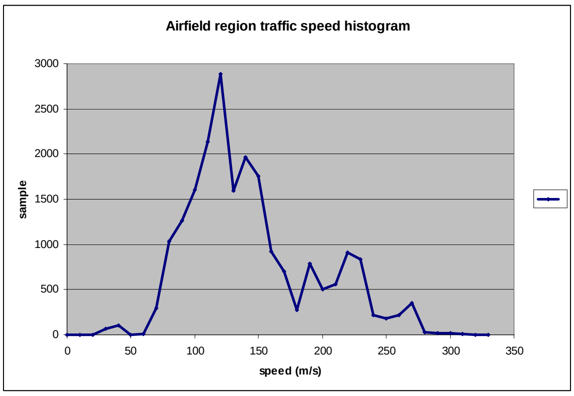

Figure 5 Speed profile for tracks in the vicinity of LGW

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 an outer range set to cover 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, see Figure 6. 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.

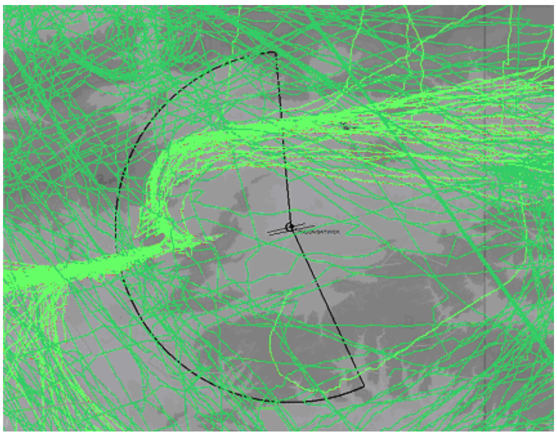

Figure 6 Area for capturing LGW departures

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 was to monitor harbours and tag all vessels leaving with origin. It was 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) provided many examples of vessel movements: AIS data provided 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 is 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 7 (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.

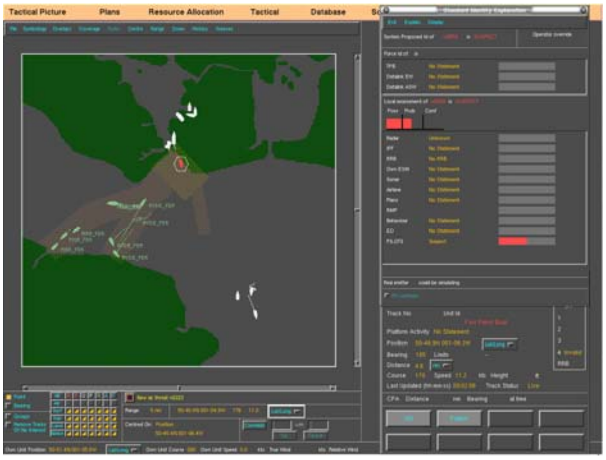

Figure 7 Vessel emerging from Portsmouth initially indicated 'suspect'

Figure 7 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.

Later the vessel emerges from the overlapping areas at the harbour mouth and now has a unique match with the Fishbourne ferry route pattern (Figure 8); 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.

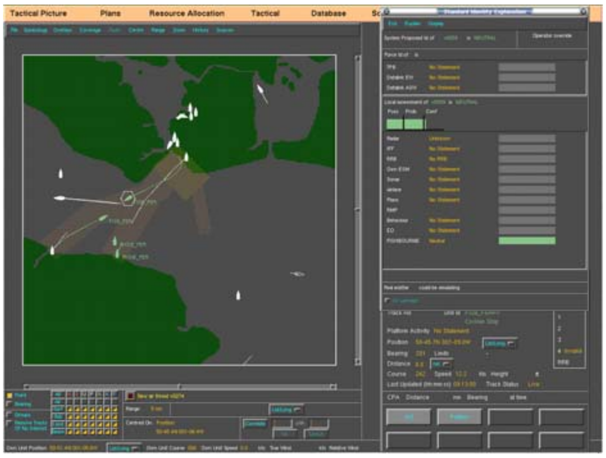

Figure 8 Vessel matching with just one ferry route

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 9 and Figure 10.

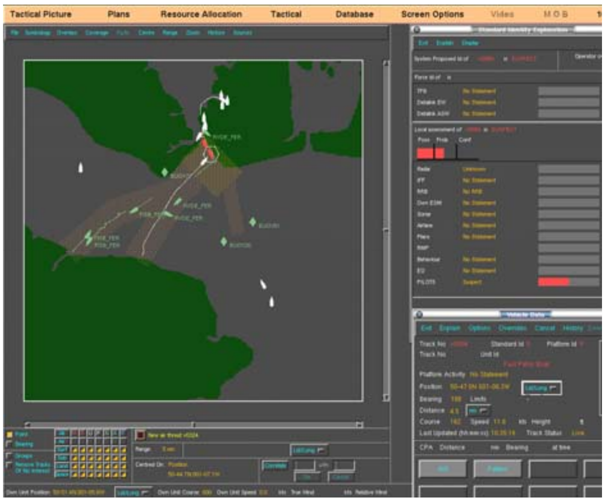

Figure 9 Two 'suspect' vessels emerge from Portsmouth harbour

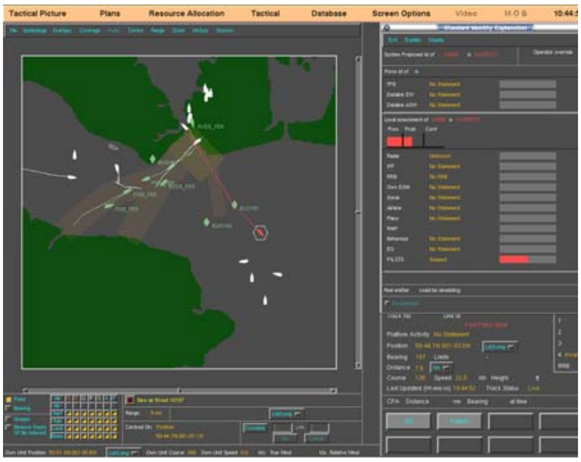

Figure 10 Vessel retains identity of 'suspect', 'fast patrol boat'

The cases above illustrate plausible operation but have 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 could 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.

Pattern matching with Synthetic data

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 11.

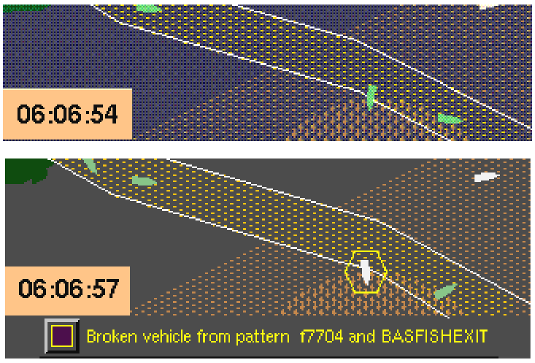

Figure 11 An example of a contact breaking from a transit behaviour

5. Emerging Issues

State of the Practice

Two factors are driving UK developments in complex combat systems. Affordability is forcing acquisition authorities to demand Open Architectures that offer the promise of new information services independent of core system suppliers. Operationally, there is a drive to exploit all available information within a deployed force, almost as a duty of care to both the forces deployed and the tax-payer.

Teams in QinetiQ are variously working on theory for future implementations in both these areas. The Level 2 fusion investigation reported in this paper has taken the robust Level 1 automated picture compilation and implemented a Level 2 fusion approach to the information accessible within a single platform, avoiding the obvious challenges of distributed data fusion across a force reported in [1].

The modular data fusion services offer a route to incrementally deliver exploitation capability consistent with a service oriented open architecture. To that end, QinetiQ has joined with BAE Insyte to provide a capability to the UK MoD to execute a sea trial to demonstrate the Maritime Exploitation of Networked Data later in 2007.

Human centred design

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 fast contacts show up in the fishing fleet as suspected FIACs in Figure 12 below.

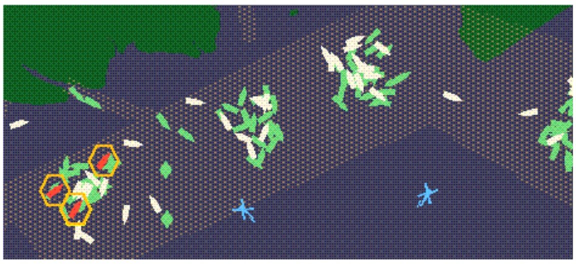

Figure 12. High speed contacts flagged within fishing vessel cluster

A combination of design and experimentation will be necessary to ensure that this Level 2 data fusion helps rather than loads users.

Knowledge Acquisition & Retention

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 13.

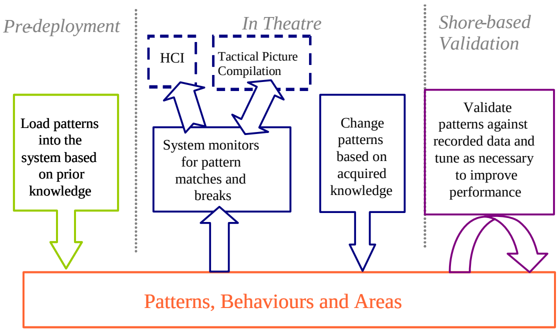

Figure 13 The life cycle of Patterns

Next Steps

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

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

6. 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 infer identity. The ASAS prototype was exercised against plausible operational examples, some taken from synthetic data representing a realistic operational context, others taken from recorded live data from a UK-based Shore Integration Facility (SIF) and using targets of opportunity.

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.

26 Mar 2007

QinetiQ/07/00887

This work has been sponsored by the Ministry of Defence, Director Equipment Capability Above Water Effects under the Combat Systems Integration Interoperability and Performance research programme.

Reference:

1. Llinas J., Bowman C., Rogova G., Steinberg A., Waltz E., & White F., "Revisiting the JDL Data Fusion Model II" *Proceedings of 7th International conference on Information Fusion, 2004.*