Defense Technical Information Center Compilation Part Notice

ADP010885

TITLE: An Overview of Information Fusion

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: New Information Processing Techniques for Military Systems [les Nouvelles techniques de traitement de l'information pour les systemes militaires]

To order the complete compilation report, use: ADA391919

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, ect. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:
ADP010865 thru ADP010894

An Overview of Information Fusion©

G D Whitaker
Defence Evaluation and Research Agency (DERA)
St. Andrew's Road
Malvern, Worcestershire, WR14 3PS
United Kingdom

Telephone: +44 (0)1684 895822
Facsimile: +44 (0)1684 894384
E-mail: D.Whitaker@signal.dera.gov.uk
WWW: http://www.dera.gov.uk

1. Abstract

This paper provides an introduction to, and overview of, the field of information fusion within a wider data and information fusion and processing context. It starts by considering the aims and objectives of research and development programmes in this area. In particular, asking what are we trying to achieve by such fusion from the end user (military commander?) point of view. The main emphasis of the paper is on military systems and reference is made to work at the United Kingdom's Defence Evaluation and Research Agency (DERA) for examples but the paper has more general relevance.

Some of the common operational and logistical difficulties associated with current information fusion systems are highlighted. In other words, "Why is making sense of data difficult?" The differences and similarities between data and information and between their fusion and processing are discussed. The rôle of information fusion systems is to address some or all of these difficulties and so provide more effective systems for a range of different applications and users. The means by which this is accomplished is then described in terms of fusing information at various levels of abstraction. Reference is made to models, architectures and frameworks that have been developed independently within the USA and the UK and that help structure and clarify the whole process.

Current research aims to further improve our capabilities in this important, force-multiplying technology. Some people and nations aspire to information dominance in modern conflicts and the same can be said for modern businesses. To achieve or even approach this goal requires a vigorous and healthy research programme. Some of the current key research activities in this area are summarised. Most of this research is targeted on specific, near-term applications. The paper concludes with a personal perspective on the main future, longer-term research challenges.

2. Introduction

The United Kingdom's (UK) Defence Evaluation and Research Agency (DERA) is an agency of the UK's Ministry of Defence (MoD). It has an annual turnover of over £1Bn and around 12,000 staff – around 75% of whom are scientists and engineers. The organisation has 15 major sites in the UK and many smaller sites both in the UK and in other countries. As such, it is one of the largest research organisations in Europe.

DERA was formed from many previous non-nuclear MoD defence research establishments and supports the MoD mainly by providing impartial advice and by undertaking research. DERA is increasingly encouraged to undertake work for other organisations on commercial terms. Such work must conform with the framework of its corporate and business plans as agreed by the Secretary of State for Defence but, even so, this extends the range of the research activities.

A survey in the late 1990s showed that there were over 100 MoD funded projects that were involved in data or information fusion. Many of these only had a small interest in this subject but there were still many that had a major data fusion activity. One might assume that with this number of projects each researching the same subject, there must be considerable overlap and therefore scope for savings measures. On investigation, however, there was actually very little duplication. Each of these projects differed from the others in some crucial way. It also illustrated the depth and breadth of DERA's research in this field.

Some of these projects were addressing land domain problems, others concentrated on maritime (surface or sub-surface), others were concerned with the air domain. This could apply to both the sensor platform and to the targets of interest. Different combinations of sensors were also being used or investigated. Some applications needed real time solutions, others could process off-line.

Paper presented at the RTO IST Symposium on "New Information Processing Techniques for Military Systems", held in Istanbul, Turkey, 9-11 October 2000, and published in RTO MP-049.

Some projects were challenged to improve detection capability, others to improve location or identification or tracking. Others were compiling whole pictures of the battlespace of interest whilst others were looking at more abstract concepts such as situation or threat assessment, decision support or even decision making. Some were concerned more with the implementation of such decisions by investigating logistics and planning.

3. Aims and Objectives

There are many military platforms and command and control (C&C) centres that need to fuse information. Initially, the aims and objectives of the different platforms within the various domains appear quite different. Submarines have the difficulties of acoustic propagation and lots of background noise. Surface ships have to deal with fast, sea skimming missiles that are difficult to detect amidst the waves. Land conflicts involve relatively large numbers of platforms each of which has considerable autonomy and is difficult to detect or predict – they are called soldiers. In the air, platforms move rapidly and can perform tight manoeuvres. In space, the main difficulties are the distances involved and observing through the atmosphere.

In most battles, there will be at least one command and control centre overseeing information across the whole battlespace but at higher levels of abstraction than most platforms. Much of the processing at this level relies on very capable neural networks colloquially referred to as human brains.

Is there anything that these domains and platforms have in common? In order to answer this question, we need to look at the underlying situation and ask why these entities wish to make sense of data. Each of these platforms, like each of us, exists within an "outside-world". We use sensors and non-sensor derived information to construct a picture of our perception of that outside world. For the moment, let us regard the data and information as entering a "black box" that contains all the processing and fusion. Coming out of the black box are generally commands and instructions that impact on the outside world. So the simple answer to the question, "Why make sense of data?" is that we wish to find out about that outside world in which we exist. This is rarely the end of the answer because, usually, we wish to influence that outside world to benefit ourselves in some way or another.

The military doctrines refer to an "OODA" loop. That is:

- they **observe** the outside world using sensors and collateral information;
- they then construct a picture representing their perception of that outside world in order to **orient** themselves within it;
- they then **decide** if and what and when and how they wish to influence that outside world; and
- finally they **act** by implementing their decision.

This act may be observed directly and / or the impact on the outside world may be observed and the loop cycles around again.

This OODA loop represents a challenge for military commanders and for the providers of technology to support or to undertake the various functions. In real life, and particularly in modern conflicts, the situation can be much more complex. There is generally an enemy or competitor as well as several neutral parties involved, each of whom is following their own OODA loop. The challenge is to make sure that your OODA loop is better and / or faster than the OODA loop of your opponent(s) or competitor(s) in order to give you the best chance of being successful in your objectives.

4. Current difficulties

There are many difficulties facing specific platforms or specific functions. There are, however, some difficulties common to many of them. One is data or information deluge. Firstly, there is a huge and increasing amount of data and information available to us. Sensors are seeing further and in more detail and in more conditions than ever before. There are also more sensors available. Non-sensor derived information such as books, doctrines, best practices, news feeds, and the internet are also growing at high rates. Secondly, researchers in the communications field have been so successful in recent years that the information is reaching us ever quicker.

In addition to the deluge challenge, information is being generated and provided by a range of diverse sources. There are significant differences in information originating from simple range and bearing sensors, compared with imaging sensors, compared with map data, compared with textual intelligence messages, compared with expert analyses, compared with information gained from past experience. Each of these provides a different challenge in its processing and then the information from the various sources needs combining or fusing.

A particularly important but often overlooked difficulty is that of specifying what is required from a system. Similarly, most providers of such systems would have difficulty in accurately assessing how effective their systems are. It can be even harder to predict the performance of sub-systems and then to aggregate these to predict the overall effectiveness for systems that are yet to be built.

An illustration of this was when the author researched and developed air defence systems. An important person from overseas was visiting the establishment. An equally important host was showing the visitor around. The visitor was shown the models and simulations of air defence systems and seemed suitably impressed. He then asked the deceptively simple question, "So how good is the UK air defence system?" The host thought for a while and then replied, "About 5." The visitor looked quizzical for a moment and then asked the supplementary question, "On what scale is that?" To which the reply was, "On any scale you choose."

Challenges that face many systems include the detection of objects or events, their location in space and / or time, the identification and recognition of those objects or events, the assessment of the situation, the making of a decision, and the planning of the implementation of that decision.

More specific challenges affecting smaller numbers of projects or domains include: low-observable targets, high noise, manoeuvres, data incest, multi-platform architectures (get data to right place at right time - push versus pull), data fusion / processing architecture map to C&C or vice-versa, conflict resolution (C&C versus own platform priorities), robust against information warfare, fusion of numeric, categorical and "soft" data, interoperability, bandwidth, confidence factors, variable latency, variable reliability, variable quality (accuracy, uncertainty, duplicates, omissions, contradictions, ambiguities).

At this point it is worth making a few observations about digital data links and about digitisation in general. Both of these are vital both in their operational significance and in the work being undertaken to develop or further improve their capabilities. They are not, however, the total solution. In particular, conformity with the data link protocols does not guarantee that platforms can meaningfully interoperate. Similarly, just because data are stored digitally, does not mean that they have been fused or that they are accurate. Whenever data are being transferred, there will always be some latency or time delay. Bandwidth is not, and never will be, infinite and so it will not be possible to transfer all the data that are available. Even if the latency were zero and the bandwidth were infinite, there would still be the issue of the data deluge and the need to extract the **relevant** data from that deluge. Hence the observation that, even with the improvements that data links and digitisation provide, data and information fusion, processing and management algorithms will be essential.

To summarise the requirement: military commanders (like everyone else) need **appropriate** information on which to base their command decisions. This information must be necessary, in that unnecessary information just adds to the confusion. It must be sufficient or the commander may be missing the crucial fact that will undermine his or her decision. It must be timely as information arriving a day late, an hour late or even a second after the decision has been made is of limited, if any, use and may even be detrimental to the future battle. Finally, it must be in a suitable format. If the information is being received by a computer and it is not in a suitable format then the receiving computer will, probably, ignore it – or crash. If people are receiving the information, then they may be tired, hungry, frightened, and uncomfortable (especially if they are in a battle where they are being fired upon). Military commanders under such stress may politely (!) suggest that such information is reformat so that they can more readily absorb it.

Data and information are provided by sensors (such as radar, sonar, optical and infra-red devices, aerial and satellite photographs) and other sources (such as intelligence, open-source, background knowledge). Modern systems outstrip the abilities of human operators, even well trained expert military staff, to absorb, process, and make all the correct inferences based on that data. This is especially the case when the information is "difficult" in some way. For instance when it is incomplete, uncertain, ambiguous or contentious. Hence there is a growing need for automated support.

5. Rôle of data and information fusion

Many people spend inordinate amounts of time discussing possible definitions of data and information fusion. For the purposes of this paper, the following, quite open, definition will be used:

***Data fusion** is the combining of data.*

Strictly speaking, this is all that data fusion is. This alone, however, does not justify their embarking on lengthy programmes attempting to fuse their data. Such programmes might be fun and may provide considerable interest, enjoyment and satisfaction for those undertaking the work. Given that funds and time are generally limited, however, one would hope that there was a purpose or objective to justify the development of this fusion capability.

This objective is usually improved performance in some way. For instance, increased accuracy, increased robustness, decreased time, increased resilience to counter-measures, less false alarms, or more effective decisions.

Consequently, a more realistic definition of data fusion is:

***Data fusion** is the combining of data ….*
….. to achieve an objective.

From this open definition it can be seen that data can be combined or fused in many different ways. It can be fused:

- **Within a single sensor at each moment in time**. For example, within an imaging sensor, adjacent pixels can be combined to form elements of the total scene. Another example is a radar that generates a range and a bearing that can be combined to deduce a geographical location.

- **Over time.** A series of observations can be used to predict a target's movements or that an engine is about to fail.
- **Among similar sensors.** A large part of the battlespace can be covered with a series of adjoining or overlapping sensors. For instance, siting radars along a coastline to detect incoming aircraft or arrays of sonobuoys attempting to detect submarines.
- **Among disparate sensors.** Using a sensor with a wide field of view to cue a more focused or accurate sensor. Exploiting complementary features of, say, infra-red and visible light to improve the identification of targets.
- **Sensors with non-sensors.** Data from battlefield sensors combined with a knowledge of doctrine may help assess enemy threats. Intelligence reports could cue local sensors to search for targets in likely regions or at specific times. News feeds and the internet can also provide a wealth of information.
- **At various levels of abstraction.** This will be expanded upon later but, for now, raw sensor data can be fused as can probabilities or decisions.

Finally,

Information *is data that is relevant to this application.*

Thus, one application's or one person's **information** is merely irrelevant **data** to another application or person. This definition also implies that there is much overlap between the underlying techniques for **data fusion** and those for **information fusion**. Consequently, the definition and expansion of **data fusion** can be repeated for **information fusion**.

To determine the rôle of data and information fusion, we need to return to the "black box" introduced in section 3. Information about the outside world was taken as input, and commands and instructions were issued as output. Now we consider the content of that black box. First, however, it should be observed that this black box may contain people and / or it may contain computers. Ideally it will contain the most effective mixture of people and computers for the particular application being undertaken. In many effective systems there is good synergy between humans and machines.

Our interest is not in whether there are people or computers but in the key components of the black box. What are the sub-functions and how do they relate to each other? It should be noted that, theoretically, such a black box need not have pre-determined sub-components. Research into self-organising systems suggest that appropriate structures can be discovered or emerge automatically without external supervision (Reference [1]).

Several groups have considered this question. Perhaps most notable is the United States of America (USA) Joint Directors of Laboratories (JDL) with their data fusion model. This has evolved over the years and an interpretation of a recent version can be seen in Figure 1.

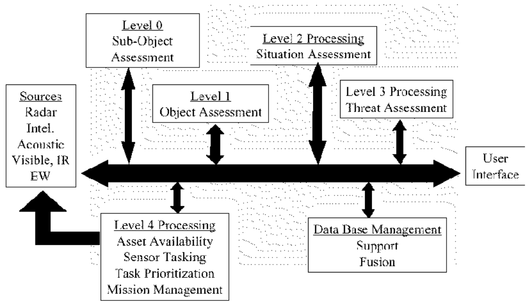

Figure 1 - USA JDL Data Fusion Model

Developed at around the same time as the JDL model but updated for this paper was one from the UK. This can be seen in Figure 2.

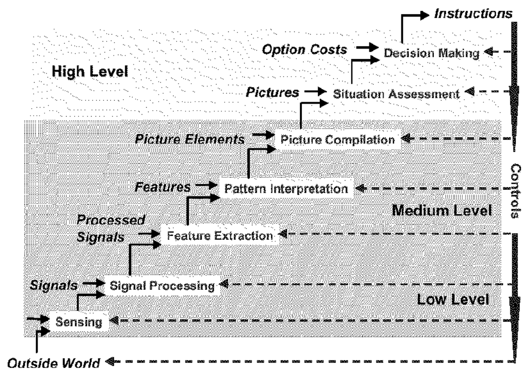

Figure 2 – A UK Data Fusion Model

Both models serve the excellent purpose of clarifying understanding and facilitating communications in this complex technical field. The USA JDL model better covers the context within which data fusion sits. For example, the important aspects of data base management and user interfaces appear explicitly. On the other hand, even in comparison with this later USA JDL version, the UK model is more explicit as to the contents of the "black box".

It may be helpful to consider the meaning of these key components:

- **Signal processing.** The initial analysis of the raw sensor data. Usually provided as a part of the sensor system and relying heavily on a good understanding of the sensor and the physics of the environment in which it operates.
- **Data feature extraction.** The extraction or selection of significant attributes from the processed data. It is important to determine the most useful features that support subsequent processing and fusion whilst ignoring irrelevant or less relevant aspects of the data.
- **Pattern processing.** Analysis of the features to estimate the elements of the total picture. These elements could be individual targets.
- **Picture compilation.** The combination of all of the picture elements into a total picture. That is, a local perception of what is happening in the outside world.
- **Situation assessment.** The determination of what the picture means from the viewpoint of interest; the determination of possible futures based on options available to each of the participants; and the evaluation of those possible futures from the viewpoint of interest.
- **Decision making.** Selecting the action such that the resulting sub-set of possible futures is desirable. The entity making the decision has a number of options that have been identified during *situation assessment*. One of those options needs to be selected (doing nothing is an option) and this should be chosen such that the future is more desirable than that resulting from any of the other options. This could be the future in which the entity is in the most favoured position. It could, however, be a good option that can be identified quickly - even if it may not be the best. This might keep up the battle tempo and get inside the enemy's OODA loop. It could be a good option that is robust to uncertainty in the picture or uncertainty in the assessment or to countermeasures. The choice will depend on the strategy, terms of engagement etc. being followed by the entity making the decision.

It should also be noted that higher levels can often beneficially influence lower levels. A simple example is the deployment of a sensor to provide further valuable information.

These components represent increasing levels of abstraction. Levels of abstraction also exist within command and control (C&C) systems. That is not to say that each and every level identified here can be found in C&C systems or that the levels can be separated out in this way. By considering such hierarchies, however, sensible questions can be posed as to where fusion should take place in order to maximise the decision making capability of the total system whilst taking account of practical constraints. Fusing at a low level provides the best prospect for optimum use of the data but requires higher bandwidths so may be limited to fusion onboard a single platform. Fusing at higher levels reduces the bandwidth requirements, but relies critically on the accuracy of information provided by the lower levels. Intermediate levels can offer compromises on these extremes.

C&C systems often have hierarchies within them. For example, a platform will, generally, have a number of sensors and data links providing input that is to be combined with on-board databases and the like. Typically there are several platforms engaged in any operation. These platforms are controlled by, for instance, an airborne C&C (such as AWACS – Airborne Warning and Control System) or a ground C&C (such as UKADGE – UK Air Defence Ground Environment). These in turn provide information and receive commands from a higher level, such as a Combined Air Operations Centre, CAOC. Thus, this example C&C system has multiple levels. The sensors will be operating on, for instance, range and bearing or pixels. This level of detail will rarely be of interest to the commander based in the AWACS who will be more concerned with the locations and trajectories of hostile targets. Even this may need to be abstracted to, for instance, the numbers of aircraft on combat air patrol in certain regions of the battlespace. The commander at the CAOC will, typically, find this too detailed and be more interested in approximate sizes of enemy forces, their likely targets, and the status of defence capabilities that could be brought to bear on this threat.

Another perspective on these hierarchies can be found in multi-platform data fusion. A package of aircraft may consist of a main attack squadron with a flight of helicopters dealing with ground forces and back-up fighters in the rear. Each of these groups will be exchanging information internally and the groups within the package also need to share summary information with the other groups. The whole package may be in communications with air and / or ground command and control stations who, in turn, will be relaying information with headquarters.

Even within one layer in this example, there are hierarchies. A single air platform will have a variety of sensor systems (visual, radar, electro-optic, identification-friend-or-foe – IFF etc.), often more than one communications system (Link 16, SATCOM, on-board communications etc), several weapons systems, tactical and geographic databases, intelligence sources etc. Consequently, each of these systems should be regarded as a system of systems in its own right.

Now consider the fusion of information at some of the levels within such a hierarchy. Firstly, fusion at the sub-system level – for example a ground based radar. The radar will **observe** by making detections that can be confirmed by temporal fusion of data and by combining data from different modes of operation. The output from the radar may well be tracks. The radar data processing system may then **orient** such data into, for instance, latitude and longitude. The system will then **decide** if mode changes are required to improve this local picture. If a change is desirable, then the system will **act** on that decision and implement the change.

Secondly, fusion at the platform level – for instance a ship or fixed-wing aircraft. Such platforms will **observe** by associating the inputs that they receive from their various sensors and communications links to form tracks. The platform system will then **orient** these tracks by fusing them and perhaps combining them with maps to form a local recognised air or sea picture. A **decision** will then be made on any action to take such as moving the platform to gain further information or to avoid a threat or to switch-on or redirect a sensor. The platform with then **act** by implementing these decisions.

Finally, fusion at the command and control centre (C&C). C&C will **observe** the inputs from its own controlled sensors and platforms. It will then **orient** by fusing this input and forming a wider picture. It will **decide** on what actions to take, and then plans will be formulated and taskings prepared. The **act** is to issue these taskings as commands to the sensors and controlled forces.

In each case and at each level of these hierarchies, there is an OODA (**o**bserve, **o**rient, **d**ecide, **a**ct) loop being enacted. Looking for levels of abstraction and OODA loops can help with the understanding of the processes. This in turn can lead to more effective command and control systems.

6. Future challenges

What of the future - where will the challenges lie? In the near term future, the challenges lie in addressing and solving the current difficulties identified earlier in section 4. That answer is a little too glib and does not aid thinking of the longer term direction. Obviously, no-one can know with certainty what new difficulties and challenges we will face. The author's personal, current opinion is as follows.

An all pervading goal is to get a better grip on specifying requirements, assessing systems and measuring effectiveness. The author makes no apologies for repeating this item from the earlier list of current difficulties. Much further progress is necessary in this vital area to set everything else in context.

More specific challenges that were not listed earlier include:

- **Improved models.** Many models make gross assumptions in order to make progress or to achieve real time performance. Such assumptions include linear motions and Gaussian noise or distributions. Where these work for specific applications, this is acceptable. Generally, however, the real world consists of non-linear systems and non-Gaussian distributions. Ask a fighter pilot if, when being targeted by a heat seeking missile, they would be happy to fly straight, level and at a constant velocity! As computers operate ever faster, the performance costs of more accurately modelling such real world characteristics become less critical.

- **Consistent uncertainty handling.** Some systems still assume there is no uncertainty in the underlying data or their resulting decisions. Denying subsequent systems this vital additional information can make the difference between a good decision and a bad decision; between surviving and being killed; being victorious and suffering defeat. Providing accurate, or at least improved, estimates of the uncertainty can, in many cases, be more important and more useful than gaining the last iota of accuracy in the decision itself. A "hard" decision, one with the uncertainty or confidence removed, can act as a veto over other systems that may be being more honest over their uncertainty and so negate many of the desired benefits from information fusion. Further advantages can be gained if each system handles uncertainty in a consistent manner. There are a number of well known methodologies for dealing with uncertainty in a sound and rigorous manner. For example, Bayes, Dempster Shafer and Fuzzy Logic.

- **Mixed type fusion.** Increasingly, modern systems are introducing fusion of information but usually at the same or a similar level of abstraction. For instance, combining arrays of pixels from imaging sensors, or range and bearings from radars, or text from intelligence reports. More of a challenge is to combine information from different levels. Techniques to do this, such as Bayesian Belief Networks or Graphical Information Models, are emerging from research laboratories but they are still limited. Particular difficulties are experienced when trying to fuse dynamic and especially fast changing data.

- **Higher levels of abstraction.** Fusion is becoming ubiquitous at lower levels of abstraction. Higher, more abstract levels are still, predominantly, left to humans. The first challenge at these higher levels is to extract the more symbolic information from the lower level data. Once this has been achieved, there still remain further difficulties. Many of the usual assumptions that we make about data no longer apply. For instance, size and proximity do not, necessarily, make sense when dealing with categorical data. Is an *armoured personnel carrier* nearer to a *truck* or nearer to a *main battle tank*? In the past, humans were being overloaded by low level data and could not see the overall picture without automated assistance. Now, people are getting a better view of the entities in the battlespace but cannot always extract and assess the situation or threat or their defensive options. There is an appropriate saying that "we cannot see the wood for the trees".

- **Resource management.** People, sensors, weapons and equipment are examples of resources that may be scarce, valuable and / or expensive. It is critical to modern, high tempo operations that such resources are well utilised to maximise the effectiveness of command and control systems. For example, sensors may be providing more data that simply re-enforce current perceptions of the outside world. It may be more effective to relocate or redirect such sensors to gain information about an area of ignorance. Similarly, improved planning of the deployment of a suite of sensors that takes into account the current state of knowledge and ignorance could reduce the time taken to detect and track a target whilst using less sensors. Other examples include the scheduling of requests for use of a resource within tight time constraints or optimising movements across route networks.
- **Frameworks and architectures.** In an ideal world, from the perspective of automating the decision making process, all data would be made available at a central point, processed and fused, and a decision made. Unfortunately, real life does not permit this. Each sensor system and each platform will undertake fusion and processing. Any information passed on to other participants will suffer a time delay. There will also be information lost due to the limited bandwidth and to fixed message types. This situation will be aggravated with each onward transmission. In a stable environment, an **architecture** can be devised that takes into account the ideal goals of the decision making process and tempers them with reality to compromise on a pragmatic system that performs effectively within the constraints. The old "cold-war" was, arguably, such an environment. Modern conflicts and operations, however, are more difficult to predict. More attention should, perhaps, therefore be given to identifying a kit of parts, or **framework**. Such a framework needs to identify useful components and suitable mechanisms for joining those parts together. When a new situation, for conflict or peace keeping or something else, occurs, the appropriate parts can be selected and connected. This can take place rapidly, safe in the knowledge that the parts will work together to form an effective (note that this is *effective* rather than *perfect*) command and control system. A consequence of this approach is that more attention will need to be paid to the interfaces between the components than hitherto. An additional benefit of this type of approach is that the individual components can be procured separately – even several competing versions of the same part could be developed if this were thought desirable – which may be attractive both financially and in reducing risk.

7. Summary

This paper has provided an overview of information fusion. It has started with the basic need for data and information fusion. There are many application specific requirements to be satisfied but the underlying need is to find out and to influence. Generic aims and objectives have been identified and discussed and some application specific requirements listed to give a flavour for the challenges facing this technology.

The rôle of data and information fusion has been described, including two models that aim to clarify thinking and aid communications for workers in this domain. The relationship between this technical hierarchy of abstractions and military command and control systems has been explored with illustrative examples at a variety of these levels. Finally the paper has indicated the author's view on the challenges that we all face in the future. Whether this is an accurate prediction or not, remains to be seen. What is clear, is that information fusion will become increasingly important in future military operations to the extent that such operations will become infeasible without increasingly automated support. Information fusion is, without doubt, a crucial force multiplier in our modern world.

8. Acknowledgements

The author would like to recognise that much of the understanding of information fusion on which this paper is based results from involvement in, and exposure to, a wide range of UK MoD funded projects. In particular, those projects funded by the Corporate Research Programme's technology group 10, research objective 4.

Many of these projects have involved staff from DERA's pattern and information processing group and it is they especially who have helped form the author's views over many years of patient explanations and shared enthusiasm for the work.

In addition, special thanks are due to Martin Ferry who contributed the original versions of some of the air warfare viewfoils on which parts of this paper are based. Similarly, many other colleagues across DERA have made contributions either directly or indirectly.

9. References

[1] Luttrell S P and Webber C J S, 1999, DERA/S&P/SPI/TR990564/1.0 (local ref. DERA/S&P/SPI/651/FUN/STIT/5_17/1.0), Using self-organising neural networks to discover structure in data.