16th ICCRTS

"Collective C2 in Multinational Civil-Military Operations"

Information Fusion for Collaborating Commanders at Different Levels

Topic(s)

Topic 9: Networks and Networking

Name of Author(s)

Martin Holmberg (martin.holmberg@fhs.se)
Swedish National Defence College
P.O. Box 27805
SE-115 93 Stockholm
Sweden

Pontus Svenson (ponsve@foi.se)
Swedish Defence Research Agency (FOI)
SE-164 90 Stockholm
Sweden

Point of Contact

Martin Holmberg

martin.holmberg@fhs.se

Abstract

This is a position paper discussing the authors' views on the role of automated information fusion in the interaction between different command levels. The purpose of the paper is to initiate a discussion on the relationship between automated fusion and the flexibility in a mission. The sharing of data/information/knowledge between commanders at different levels is a difficult task in many respects. We focus on the role of automated information fusion techniques in this frame. The paper asks two major questions: (1) Does automated fusion generate the unwanted side-effect of less flexibility? (2) How should a situational picture be represented at different command levels to promote cooperation? We also discuss some disadvantages of using traditional information fusion methods developed to handle either high or low level information. The conclusion is that information fusion techniques have to be chosen with care when making information systems that should be jointly used by commanders at different levels.

Introduction

Traditionally, military sensorsystems (consisting of one or several sensors and some signal processing/data processing system) are designed to provide information to an operator, or a local user. This operator can be trained on the system, learn to interpret rather complex outputs, and use the output for a specific task. This means that sensor systems often are designed to fulfil one (or a few) specific task for an operator with a given background knowledge, and in this context provide a robust and accurate output. Very seldom, the focus has been to design a system with a high interpretability for untrained users since that could easily affect other performance measures of the system, such as speed or accuracy.

As data from many different sources becomes available to different users at different command levels, there is also an increasing need to consider in what form the data is available. For a given (sensor) system, instead of being adapted only for use by the operator of the system, the output from the system has to be in a form that makes it useful also for other users. So, the system has to be designed to fulfil not only the needs of the operator, but also unknown users with possibly completely different requirements.

At a tactical level, the user is often interested to find primarily (near) real-time information on individual objects such as their position and class or identity. If the system does not provide information of sufficient detail, the operator knows the system well enough to look at raw data (e.g. a camera image) directly. Also, since the distance between operator and sensor system is short, the band width for transmitting the data is rarely a problem. Operators at operational or strategic level may be interested in data from the same sensor system, but they rarely have much knowledge on how the system works, and they may also be interested in data in different forms, e.g. historical data or the data related to a context. In this paper, we refer to these levels as different command levels.

Information fusion is a mathematical process of combining data from different sources (e.g. sensors, human intelligence) in order to reach a better description of a system than is possible using only a single source. This can also be made at different levels, e.g. according to the JDL model, where the lowest level deals with describing single objects, the next level sets the object(s) in a setting that can be interpreted as a situation description for the user, and a third level evaluates what threats/possibilities that exist given the information available. Also further levels can be imagined. These levels do not correspond to the different command levels, and we refer to these levels as fusion levels.

In this paper, some needs of operators working at different command levels are discussed, and the implications of these needs for fusion methods are discussed. Also, the effect of automated fusion on the flexibility of a mission is discussed. The aim of the paper is not to present new scientific results, but to start a discussion on how command and control is affected by technical systems in ways usually not foreseen.

Related work

In recent years, there has been a large amount of work on information fusion (Liggins et al 2009) and its relation to situation awareness. A specific focus has been on how to combine hard (sensor) and soft (human-generated, e.g., text) information, and there has also been a lot of work on using semantic technologies for fusion. For semi-comprehensive overviews of research in high-level information fusion, we refer the reader to the FOI reports (Hörling et al 2010, Svenson 2008, Berggren et al 2007).

Little and Rogova (2009) discuss the important problem of how to design ontologies suitable for high-level information fusion. They argue that it is necessary to distinguish between events that are spatial or temporal and present a large number of relations between them that are necessary to deal with. The work is exemplified with a formal ontology for situation and threat analysis in crisis management. It is not enough just to have ontologies, they must also be used and populated. Hull et al (2009) present a method for semi-automatically tagging messages about criminal activities. The goal is to make it easier for human operators to find information about, e.g., car thefts. Keywords along with sentence structures is used to determine what tags to use for a specific message.

Parallel analysis of text by several knowledge processes is the subject of a paper by Auger (2009), which describes a Multi-Intelligence Tool Suite. The different processes each have different knowledge bases and rules for classification of the text.

Kokar et al (2009) present a review of the situation theory due to Devlin. This is based on infons, tuples that express a relation between a number of objects (which in turn can be infons). Infons also have associated truth values and can be combined to express complex logical relations between them. A situation is said to support an infon if the infon is true in that situation. The paper goes on to introduce a semantic adapted to computer inference and

compatible with situation theory as well as Endsley's theory of situation awareness (Endsley 1997).

Crisis management is also the focus of a paper by Baumgartner et al (2010), where an ontology-based framework for increasing situation awareness is presented. They assume that each information source has its own domain ontology, and introduce mapping ontologies that translate between them. A reasoning engine is then used to determine if crisis events are happening, and the operators should be alarmed. Evaluation results from traffic control in Austria are presented. The system successfully manages to increase the awareness of the operators, but requires domain experts and creation of rules that cover all possible events. The system is not capable of discovering new events for which there have been no previous rules.

Sensor networks

Sensor networks are assembles of sensors spread in an area of interest and autonomously collaborating to perform a task defined by a user. In this work, we will deal only with networks of small/moderate size, covering up to a few square kilometres, consisting of small sensors without external power supply, and used for surveillance purpose. These networks are used for tactical needs, and the result is usually delivered to a local user. The network delivers information like position and class of objects present within the network range. However, the information from the network can be important also for users on a higher level (operational or strategic) e.g. when the presence of certain objects has strategic implications. It is therefore necessary for the network to deliver the information in a form useful for users at different levels. The information should be accessible both as "information-push" (i.e. the network autonomously delivers relevant information to the user), and as "information-pull" (i.e. the user should be able to search for relevant information in the area of interest, which can include one or several sensor networks, human intelligence, etc).

Research in the sensor network area is, in general, focused on one specific part of the system, e.g. communication, data fusion, or sensor nodes. The sensor network can for example be optimised to use little energy for communication, give precise results after data fusion, or use sensors with low power consumption. In order to make optimisations like this, several parameters in the network can be affected: Routing protocols can be changed, methods for performing distributed data fusion can be developed, or hardware parts can be replaced. Many of these changes come with a price, for example using less power consuming routing algorithms may lead to a less robust network, data fusion methods may require more communication and thus consume more energy, and sensors may become less flexible when manipulated to consume less energy. It is therefore required to combine all these aspects in order to study the sensor network as a complete system, something which rarely is done. An example of a more complete system approach is to optimise the performance of the data fusion models, given the constraint that the energy available for measurements, calculations, and (primarily) communication is restricted.

When viewing a sensor network as a part of a greater system, even more aspects become relevant. What is the information need of the (remote) users? Can this be met by the sensor network? A sensor network can only deliver the information it has been designed to extract, and if remote users require different information than a local user the network probably cannot deliver the required information. For example, if a local user requires tracking and type classification of vehicles, and a remote user needs to know the colour of all vehicles we cannot be sure that this information can easily be retrieved even though cameras in the network in theory could provide this information. On the other hand, if the network is designed to extract all information that possibly can be extracted from the sensors, the amount of data that has to be transmitted in the network will be enormous. We therefore need to consider the information needs of different types of users when developing fusion algorithms for sensor networks. But how can this be done?

We will never be able to predict the information that all potential users may require, so we cannot resort to data fusion algorithms that explicitly extract the required information in an optimal manner. Instead, we need to use algorithms that handle data in a form that easily can be related to a form that people tend to use. As an example, people tend to describe the position in terms like "close to the wall", rather than as "x=13457, y=86234". If we let the network use the latter representation, the energy consumption of the network will probably be lower during the fusion process. On the other hand, when a user requests information of the former type, the network needs to calculate that type of information directly from the raw data, it will not be readily available. In a real-time scenario, it may be difficult enough to handle this problem. If we consider the case when a user is looking for historical (from a few hours to several weeks old) data in stored files from the sensor network, we cannot assume all the raw data to still be available since that would require too much storage space. Instead, only the fused results will be available for a limited time (in some networks data can be assumed not to be stored internally, but in an external server). Again, if the fused result is in a user-adapted form, this type of search can maybe be performed. At least the probability for finding the information is higher than if the data is stored only in a computer-friendly format such as numbers with a meaning only for the network itself and the local user who knows exactly how the network is configured.

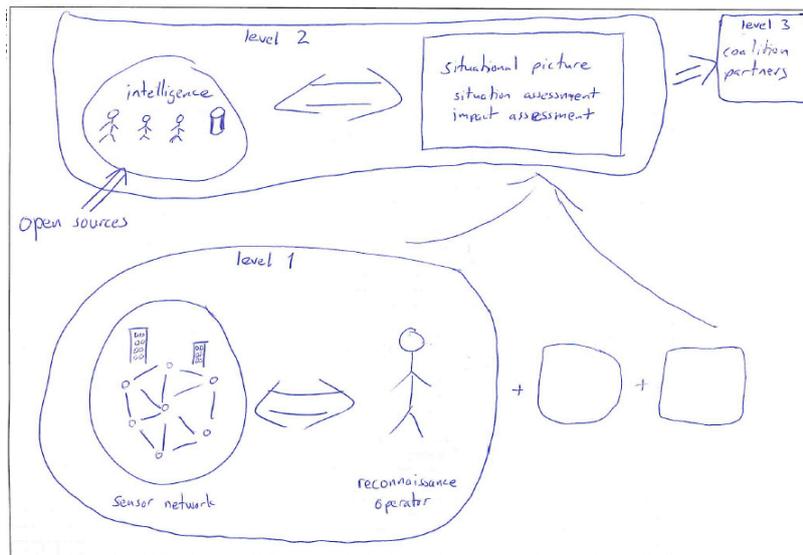**Fusing information using semantic networks**

Fig. 1. Different levels of information that need to be combined to obtain a common operational picture in the field.

In Fig. 1, fusion processes at different levels are described. The levels are introduced here only to make a distinction between the different needs and methods of different people and different sub-systems, they have no direct connection to levels used in other contexts such as command- or JDL levels. In reality there may be more levels involved, but the figure shows the principle rather than an exact architecture.

The first level shows a field operator, who has the task to perform reconnaissance in an urban environment. The operator often looks for indicators such as slogans written as graffiti on the walls, politicians visiting religious leaders, etc. If we assume that a sensor network is placed in the vicinity, the operator can tag certain objects such as a person or a car and give the task to the sensor network to track the object and gather as much information as possible. Accordingly, the sensor network may discover interesting objects and ask the operator for refinement of the information (e.g. improved classification) or alert the user when a possible indicator has been discovered. The network may thus learn also from these events. One big challenge at this level is the association problem, ie, how to be sure that the sensor network and the human refer to the same object.

The second level shows how information from several field operators, sensor networks etc. are combined to form a situational picture for the national level. As indicated above several intermediate levels may exist. The situational picture is combined with intelligence information and is processed by intelligence analysts. One of the many difficulties that arise when including sensor data in this level is that the tagging of information sources may be impractical when several fusion steps have been performed. This means that the same data may be re-used many times, and thus given too much importance (also known as data incest).

The third level in the picture shows how the national situational picture is fused with that of different coalition partners. From a fusion point of view, it would be advantageous to fuse data/information at all levels, but this is still not feasible for political reasons. In this process,

data incest is even more of a problem since the data probably is exchanged at a rather high level.

Discussion

Current decision support and situation awareness systems do not perform well in fusing pre-known intelligence or context data with sensor data. Some specialist information systems exist that are focused on intelligence and the semantic relationships between the different components of an intelligence picture, e.g. in order to tackle organised crime. These systems could possibly be extended to incorporate the data from the sensor networks and thus advance the current capabilities for integrating sensor and context data.

A key element in presenting an accurate situational picture is to perform all of the information processing in real-time, presenting information with minimum latencies and providing operators with a representation of the warfare environment as it is occurring. Related to this, time horizons are very important for decision making. A piece of fused data presented to the operator will only be relevant to a situational picture for a certain time horizon, beyond which it becomes irrelevant. To achieve both of these elements, novel approaches to both manipulating (through algorithms and a well designed System Architecture) and representing the time factor (through the latest visualisation techniques) have to be developed.

Today there exist a wide variety of different data mining algorithms that can be used to detect interesting patterns and deviations from normalcy (anomalies) in data. Of particular interest are matrix-based methods and stream-processing methods. Both of these areas will need to be further developed. An interesting data mining technique today is to use latent semantic analysis (originally from natural language processing) and topic models (latent Dirichlet allocation) to find explanations of observations and connections between observations. One challenge is to extend these methods to handle heterogeneous data from a variety of sources, including both sensor and text-based information.

Current use of data mining on sensor network data is mostly done off-line and at a low level of decision-making. Data mining of other kinds of data is often used at higher-levels of decision-making. To solve the applications above, we must bridge the gap between these two uses of data mining. Visualisation of the data mining results will also need to be addressed. Different levels of command will require different visualisations.

Current research on intelligent user interfaces provides methods to tailor the information presentation to the momentary context-of-use. In a similar way, methods are being developed to attune the allocation of tasks, the provision of cognitive support and the presentation of information to the momentary operator capacities. Integrated user interfaces can help to overcome the cognitive limitations of operators (e.g., cognitive biases) that are partly caused by dynamic situational conditions. Another very important issue is the display and manipulation of uncertain information. There are a number of different methodologies for

managing uncertainty, and more research is needed on how to construct computer systems that enables humans to understand and interact with these uncertainty structures.

Research on Computer-Supported Cooperative Work has provided diverse solutions for "collaboration at a distance", among other things to construct and maintain shared situation awareness. Questions remain on (1) how to adequately deal with different abstraction levels in the decision-making processes, and (2) how to support the maintenance of an adequate trust level in these processes, both for the technological and colleagues' information processing capacities. User interfaces must be developed to provide solutions for these two questions by supplying information access at different abstraction levels and supporting the assessment of information on aspects of ambiguity, topicality, noise and reliability. Mixed-initiative dialogue information can be added, annotated and consulted to improve the shared situation awareness.

The concepts situation picture and situation information are often used in non-ideal ways. While it is necessary to have access to all available information, it is not enough to display this information in a "situation picture", most often a map with icons representing military units on it. The situation information needs to be processed and abstracted so that the information is presented to different users in the way that is correct for that particular user. In any given situation, there is not *one* correct situation picture that can be displayed to all users and will lead them to make the correct decisions. Rather, the available situation information must be processed and filtered defending on the user's specific needs to display different views of the available situation information. These different views must of course be consistent with each other – but they are not the same and it will certainly be the case that information that is vital to one user it not even displayed to another. For instance, a platoon commander who is executing a humanitarian escort mission in neutral territory is much more interested in finding information about available petrol supplies than in medical supplies. Of course, if circumstances change, the information needs of the commander also change – if there is new information that the cease-fire has been broken, then it quickly becomes important to display the information about medical supplies to the user.

What specific information needs a user has will depend on many factors. In addition to the mission that the user is performing and their role, the situation itself also affects the needs, as exemplified above with the supplies case. We can term the different ways of presenting the situation information as showing different abstractions of it to the user. Returning to the traditional situation picture with enemy units on a map, it is for instance obvious that whereas the platoon commander needs to be shown individual enemy tanks, the battalion commander should only see enemy companies – the information displayed to them is abstracted differently.

A related issue arises when one considers the difference between data-driven and model-driven analysis. Models must be constructed beforehand, and can only detect events which humans have thought of as plausible. Data-driven analysis gives the possibility to detect unknown situations. But even data-driven analysis and the use of for example machine

learning to learn a "normal situational picture" in order to do anomaly detection are susceptible to problems when unknown, unexpected things happen since it is very difficult to define what is a normal picture. This is easier for low-level problems – for instance, to monitor a corridor and detect abnormal movement patterns (e.g., somebody running) is easy, since there are standard ways of analyzing the video. For high-level situations, the variation of what is normal is much greater, and defining a "normal situational picture" is thus much more difficult.

Since different information has to be presented to the users when collaborating between different levels, the information has to be elaborated (or fused) automatically or manually before being passed to other levels. Automatic fusion is advantageous in many cases, but may lead to a loss of flexibility since all models, and thus interpretations, are set beforehand. Consider a simple example: Several data sources indicate incoming threats in different areas that need to be handled by a low-level commander. The information is uncertain, and the commander needs to find out which data sources that are correct. He/she is also able to take appropriate steps to protect the own forces, the assets, and the civilians under protection. As long as the situation is well-known, or at least has been thought of before the mission, automatic routines may be able to handle the information and present some form of a situational picture to higher level command. However, if something completely unexpected is about to happen, automatic fusion routines and fused situational pictures cannot be trusted since the fusion routines only handles the types of data they have been designed to handle in a way (or in several ways) that have been thought of beforehand. There is no routine which comes up with crazy ideas that are very improbable to occur. On the other hand, humans handling data may do just that, think of very improbable events and examine their plausibility. Furthermore, it may not be clear how small state changes are handled by the automated fusion methods, maybe there is a risk that a small state change is lost in the uncertainty handling of the method. A human operator is more suited to direct the attention to the points of interest in each moment, and it is possible that humans therefore are better at discovering changes in the situation. The capability to imagine improbable events, and to discover small but important changes are two important sources for flexibility. So, working with automatic or semi-automatic methods to understand what is going on may help in pre-defined situations, but may also lead to an un-flexible way of interpreting the data in unexpected events. The low level commander and his/her staff can utilize all available information and respond quickly and thus be considered to be flexible, also when reality differs from the situational picture. For a high level commander to be flexible, there are two choices: Either to examine all data (which would require too much time), or to give the low level commander complete freedom of action. The latter may be difficult if the (automatically generated) situational picture indicates that the low level commander is making the wrong decisions. The flexibility of each commander will depend on the information need on his/her command level. It is therefore our belief that the exchange of information between different command levels cannot be automated to a great extent in order to retain the flexibility in action.

Finally, tools are needed to support operators and commanders in collaborative decision making. Distributed collaboration yields challenges not present in co-located teams, i.e., less

frequent communication, non-verbal communication is not possible; communication often takes place in a formal setting, while informal communication has proven to be important for sharing and synchronizing ideas and developing common ground. We therefore have to involve aspects such as:

- Information management. Prevent information overload problems, filter irrelevant details
- Team situational awareness. Team member must have a common understanding of the environment
- Team awareness and shared intent. Know each other's tasks, capabilities and workload. Synchronise each other's goals and give guidance to decisions that must be made

Conclusion

In the future, the need for including data/information from other sources than what has traditionally been used will increase. This means that technical systems supporting commanders at low levels, e.g. sensor networks, will benefit from using information from traditional high level information sources, such as intelligence reports regarding the threat level in the area. There will also a be bigger need for information exchange regarding the operation picture between different coalition partners, something that for political reasons is made at a rather high information level causing problems such as data incest. There exist some methods to handle some of these problems, but this is also a field very much open for new research.

Furthermore, data fusion techniques rely to a great extent on methods using pre-defined rules. This is very good in situations where a predictable behavior of the operator is required. However, there is also a risk that the flexibility of the operator is reduced due to the fact that the automated fusion routines are always based on the same assumptions even though the operational environment is constantly changing.

Decision support systems, based on e.g. information fusion, have many benefits and will be used extensively in the future due to the many advantages. However, being aware of possible unwanted consequences of automated routines makes it possible for users to make a conscious choice whether the desired benefits are greater than the risks for a given task.

We would like to challenge the participants to discuss these issues regarding the interaction of information fusion techniques and the operators, since we believe that there will be a large demand for this type of knowledge in the future.

Acknowledgments

References

Pontus Hörling, Marianela Garcia-Lozano, Tove Gustavi, Fredrik Johansson, Ronnie Johansson, Christian Mårtenson (2010), Omvärldsanalys informationsfusion 2009-2019, FOI report FOI-R—3066—SE (2010), available at http://www.foi.se/upload/projects/fusion/FOI-R--3066--SE.pdf

Pontus Svenson (2008). Omvärldsanalys Informationsfusion 2007 – 2008, FOI report FOI-R—2536--SE (2008). Available at http://www.foi.se/upload/projects/fusion/FOI-R--2536--SE.pdf

Peter Berggren, Pontus Hörling, Christian Mårtenson, Johan Schubert, Robert Suzic and Pontus Svenson (2007). Viktig informationsfusionsforskning i omvärlden under 2006., FOI report FOI-R-2252—SE (2007) .Available at http://www.foi.se/upload/projects/fusion/FOI-R--2252--SE.pdf

Martin E. Liggins, David Hall, James Llinas (eds) (2009), Handbook of multisensor data fusion;:theory and practice, 2d ed, CRC / Taylor & Francis (2009)

E G Little och G L Rogova (2009). Designing ontologies for higher level fusion. Information Fusion, Vol. 10, Nr. 1, Januari 2009, s. 70-82.

R D Hull, L Lashine och D Jenkins (2009). Semantic enrichment of multi-intelligence data within a net-centric environment. Proc. SPIE, Vol. 7345, 734509 (2009).

A Auger (2009). Acquisition and Exploitation of Knowledge for Defence and Security. NATO RTO IST-087 Symposium on Information Management – Exploitation, Försvarshögskolan, Stockholm, Oktober 2009.

M M. Kokar, C J Matheus och K Baclawski (2009). Ontology-based situation awareness. Information Fusion Volume 10, Issue 1, January 2009, Pages 83-98.

N Baumgartner, W Gottesheim, S Mitsch, W Retschitzegger och W Schwinger (2010). BeAware!—Situation awareness, the ontology-driven way. Data & Knowledge Engineering, Vol. 69, Nr. 11, November 2010, s. 1181-1193.

Endsley, M.R. (1997). The role of situation awareness in naturalistic decision making. In Zsambok, C.E. & G. Klein (Eds.), Naturalistic decision making (pp. 269–283). Mahwah, NJ: LEA.