14th ICCRTS: "C2 and Agility"

Development of Computerized Support Tools for Intelligence Work

Paper ID: 048
Topic: C2 Architectures and Technologies

Joel Brynielsson, Andreas Horndahl, Lisa Kaati,
Christian Mårtenson, and Pontus Svenson
Swedish Defence Research Agency
SE-164 90 Stockholm, Sweden
E-mail: `firstname.lastname@foi.se`

Abstract

In the tasks facing the armed forces today there is a need for new and improved intelligence analysis tools. The opponents no longer follow strict doctrines that determine their behavior and force-composition. Several different opposing groups must be taken into account, some of which will appear to act friendly towards us. In this paper, we describe a vision for how various information fusion tools can be used to help intelligence analysts and decision-makers achieve situation awareness. We consider intelligence work and propose an analyst-centric toolbox aiming to help analysts involved in the intelligence production process to prepare suitable reports. Intelligence analysts are overwhelmed by information, both in the form of sensory data, text stemming from human observations and other sources. We describe parts of the intelligence process and touch upon the subject of what parts can and cannot be automated. The toolbox is outlined by describing a number of possible tools, e.g., semantic information tagging, a threat model construction assistant, a situation picture construction assistant, social network visualization, a game-theoretic reasoning engine, etc. Some of the tools described have been implemented as concept prototypes whereas others are the subject of ongoing research.

1 Introduction

Intelligence analysts of today are weighed down by information that they must take into account when producing their analyses and assessments. They must produce far more content today than previously while at the same time taking into account material from a wide variety of sources, ranging from IMINT (image intelligence) to text articles in newspapers. Both the pace of the information push to the analysts and the information pull by the commanders and decision-makers who need output from the analysts has also increased: the timescales involved at operative and tactic levels are shorter compared to the days of the cold war. The types of conflicts that we are involved in have also changed: in peace-keeping and peace-enforcing missions, we are faced with a multitude of actors who are of different types and it is not always clear whether they should be regarded as friends, foes or neutrals in a given situation. When analyzing, for instance, a report on a confrontation between two clan leaders, it is important not only to know the interests, capabilities and motivations of these two, but also those of the other actors in the area who are connected to the involved parties. The goals of network-based defense and network-based intelligence will further exacerbate these problems: more sensors mean more data for the analysts to consider, and also more decisions to be made regarding where to put the sensors. The need for continuous action can make the traditional intelligence cycle obsolete. Instead, it is necessary to make all stakeholders, including customers, involved in the intelligence process, enabling a creative collaboration process where several analysts may contribute in parallel to a continuously refined shared picture of the target, see, e.g., [8] who proposes a "target-centric" intelligence cycle.

In order to meet these requirements, intelligence analysts need better information handling tools and concepts. Fusion tools [44] can be an important help, enabling, for example, automatic clustering of similar reports and using aggregation methods to produce meaningful labels on the information displayed to them. Techniques from natural language processing and text mining can be used to fuse information in different languages and to produce summaries of vast amounts of textual data.

Complete automation of the intelligence production process, however, is neither possible nor desirable. Hence, the task for fusion researchers is both to build automatic tools that process information and to build tools that help humans to do further processing: ultimately, fusion is a process involving humans. Situation awareness is created in human minds, not in machines. In this paper, we present some ideas on how computer tools for intelligence analysis can be created, i.e., tools that can be used at the discretion of a human analyst to produce intelligence products in alternative ways. What is needed, rather than constructing one single tool, is a "toolbox" containing several different tools, each helping the analyst with one specific task.

It is important that the different tools in the toolbox that we suggest can be used together; it is necessary for them both to share data format and "look and feel" in their user interfaces. We stress that the main contribution of our paper does not rely on describing any one tool in detail, but rather in presenting a smorgasbord of connected tools that aim at helping the analysts. The state of the presented tools varies considerably. In one case, we have performed extensive user-experiments with the tool, while in others we have merely a concept prototype and in yet others only an idea for the tool. We would also like to stress that the tools we are proposing are meant as aids to the users, not replacements. While the construction of intelligent computer programs that are able to "think" is a worthy goal of academic research in artificial intelligence, we are far from achieving this. And for the delicate matter of analyzing intelligence, it is not clear whether automation would be desirable even if it was possible.

The paper is outlined as follows. We first present background information on the new requirements entailed by the new kind of opponents we are facing, followed by introductory sections on information fusion and intelligence analysis. A possible architecture framework for intelligence tools is then briefly described, followed by the main section of our paper, which describes a number of analysis tools. Finally, a discussion wraps up the paper.

2 Background

In recent years, the challenges facing the Swedish Armed Forces have changed. Instead of a situation where we face one well-known opponent and the only other actors on the battle field are either Swedish civilians or Swedish soldiers, the Swedish Armed Forces are now subjected to situations where there are multiple, possibly warring, factions that oppose each other. There are also civilians present whom we must protect from harm. In addition to the inhabitants of the country where the operations are taking place, there are also NGO's (non-governmental organizations) there to provide humanitarian aid or protect human rights. The attitude towards the blue forces of these different actors can range from openly hostile over neutral to friendly, and might change on a day-to-day or week-to-week basis. The complex situations that we are now facing bring about that we must develop new methods for information fusion that are adapted to deal with the kind of information that will be available. However, we have considerably less background information than before. We can no longer rely on doctrines for how the enemy should behave when, e.g., attacking a hill. We also cannot rely on signature libraries for object classification in sensors. Our opponents will use the same equipment as the journalists and NGO representatives that are also present in the area. This calls for new ways of collecting information. Sensors of various kinds will always be important, but the most important information collecting resource in OOTW (operations other than war) will be human observers.

Soldiers who are out in the field will have access to PDA's and other equipment that will allow them to send text messages as reports. As will be detailed later, they should also be able to tag the reports with meta-information. Some of this meta-information will be directly usable in the analysis methods described below, while others will need to be analyzed and combined.

In all OOTW, civilian cooperation is important. The decision support systems thus need to be adapted to this. For example, functionality for importing material from NGO knowledge bases (KB) as well as from media reports and from the web to our own KB needs to be included in the decision support system. See [4] for an example solution for importing information from one KB into another.

The fact that the opponent does not follow known doctrine does not mean that it is impossible to analyze them or to make models that could be useful for predicting their future behavior. We believe that robust Bayesian statistics will be an essential tool for doing such analyses. Robust Bayesian analysis distinguishes itself from ordinary Bayesian analysis in that it uses an ensemble of priors and likelihoods instead of just one [5]. This means that it is possible to form hypotheses regarding what possible priors and likelihoods might be useful and then perform calculations that will give sound and robust boundaries for the expected behavior of the opponents.

As can be seen from the above, providing the force with adequate decision support tools requires the development of several different sets of tools. The information collected by sensors and human intelligence (HUMINT) needs to be analyzed and fused to create situation and impact assessment.

Historically, most technical decision support research has focused on quickly processing information from sensors to construct situation pictures and impact assessments. However, the increasing use of computers both by blue force soldiers and by people in operational areas mean that more and more text documents will be produced that contain information that is useful for achieving situational awareness. Thus, there is a need for building decision support systems that integrate (fuse) handling of textual reports with the handling of structured reports from sensors.

Planning the best use of resources in order to achieve the goals set out for the mission will also be more difficult in OOTW situations. The added complexities of analyzing several antagonists as well as different groups of civilians are one reason for this difficulty. It is likely that the use of sophisticated methods for stochastic simulation will be necessary to take account of the uncertainties present in the behavior of these actors. Another possible approach would be to combine game theory [24] with robust Bayesian analysis to handle the uncertainties.

3 Information Fusion

Information fusion [18] is a necessary and vital component of future information-sharing systems in support of command and control. It is an area that has developed considerably during the last years and continues to do so. At the start, the focus of most information fusion research was on detecting and tracking moving objects, most often in the air or on the sea. Soon, it became interesting to look not only at single objects but also at what relationships are present between the observed objects. Determining such relations is a part of situation assessment, which can be defined as providing a user with the necessary information in order for them to get a good situational awareness of the current situation. In addition to knowing where an object is, it is also interesting to determine where it might be going and what possible threat it might pose to us in the future. Determining this is called threat or impact assessment. In addition to the study of single objects, we can also look at how the relations that we determined between observed objects, terrain, our resources and other things evolve in the future. Analyzing the future behavior of the surrounding systems is also called impact assessment. It is natural to think about short-term impact assessment as being part of situation assessment. It is impossible to have a good situational awareness without also having some belief about how the observed relations will change in the near future. In impact assessment, however, we might also be interested in analyzing hypotheses about the long-term behavior of the world. This is a considerably more difficult problem.

In recent years, the information fusion research area has expanded into using information from non-physical sensors (such as text) and to answering questions that are not related to observable physical objects [13, 28]. Instead of analyzing a group of tanks moving towards us on the battle-field, the information fusion system could be used to estimate the probability of a terrorist attack happening.

As indicated above, on higher levels information fusion concerns the actual production of basic data for situation awareness and, subsequently, the prediction of future events. Hence, it follows that higher level information fusion ought to be considered an integral basic condition for the intelligence analysis process.

Intelligence analysts must be able to manipulate uncertain data, to formulate hypotheses and to test them. Languages and systems for intelligence information management need the ability to express and reason with incomplete and uncertain information. In a law enforcement intelligence system, how should the uncertainty or likelihood of each recorded entity be estimated and entered into the system? How should quantitative reasoning about alternative courses of action, or possible consequences of critical decisions be carried out in law enforcement applications? How sensitive to variations in the input parameters are the conclusions reached?

Providing such tools and methods for handling uncertain data is one important reason for why information fusion is important for intelligence work.

4 Intelligence Analysis

The work described in this paper emphasizes intelligence analysis, i.e., the work of intelligence analysts which can best be described as the art of creating useful intelligence products. A requirement for high-quality intelligence analysis is that the data to analyze is of sufficient quality. When faced with ever-increasing amounts of information to analyze and shorter and shorter time to do the analysis, it is vital that the quality of the produced intelligence does not decrease. Commanders and decision-makers must still be able to rely on the results of the process. It is therefore important that intelligence products come with a marking that states its level of quality and confidence in the presented results, as discussed in, e.g., [3].[1] In addition to such metadata, intelligence data also needs to be semantically tagged to enable quicker searching of information. In order to produce a fusion result, the fuser (whether it is a machine or a human) must first find all relevant information. For sensor data that is about a given object, this is the association or clustering problem. For text data, text clustering could be used. When combining sensory and textual data, it is possible to use semantic queries, but this requires that all data is semantically tagged—a process that needs to be automated as much as possible.

The ideas for tools described in this paper are envisioned to be used together, and they will also be combined with other tools that are currently developed at the Swedish Defence Research Agency (FOI). It is, however, also important that each tool in the intelligence analysis toolbox is useable (and use-worthy) by itself. It will be impossible to postulate beforehand what combinations of tools that are useful for an analyst working on a specific case. Instead, the analysts must be able to choose for themselves what tools to use. This thinking is in line with the emphasis on a service-oriented architecture which is affecting the current development of command and control and intelligence analysis systems in Sweden.

In addition to the technical research described herein, we also need to develop processes and methods for how humans should use the tools. These processes need to be adapted to the way that human analysts work. What parts of the toolbox that will actually prove to be useful in this process cannot be determined until the tools have been developed further and subjected to user experiments.

The intelligence analysis process can generally be divided into four phases: direction, search (or collection), analysis and presentation/dissemination. The phases are generally connected to form a loop in the so-called intelligence cycle; it is, however, important to realize that intelligence analysis is not a linear process, but involves jumping back-and-forth between the different phases. During the search phase, information is gathered from different kinds of sources, such as databases, sensors, and newspapers. In order to reduce the time spent in the search phase, a new phase dealing with information input and structuring of information can be introduced. Planning is the phase where the analyst and their customers prioritize among information requests and determine what resources (both collection and analysis resources) should be allocated to them; it corresponds loosely to the adaptation or sensor management level of information fusion.

If information is structured using semantic techniques, it might be possible to turn the solid curve in Figure 1 into the dotted one. Decreasing the amount of time spent on searching for information enables the analyst to spend more time on the later phases which, typically, involve more creative thinking and analysis. Such benefits could prove to be especially useful in time-critical situations, when it is of utmost importance to gain situation awareness in a short period of time. It is in this structuring phase of the intelligence process that the semantic tagging tool will fit. We believe that some parts of this process can be automated

[1] An interesting research challenge would be to try to change the process of disseminating intelligence products so that it could be possible to produce a "first estimate" of an analysis given to the intelligence consumers and used in, e.g., preliminary planning while waiting for the final analysis result to appear. This would be similar to how anytime algorithms [10] are used to produce better and better results while at the same time being able to always provide their current best estimate of the final answer.

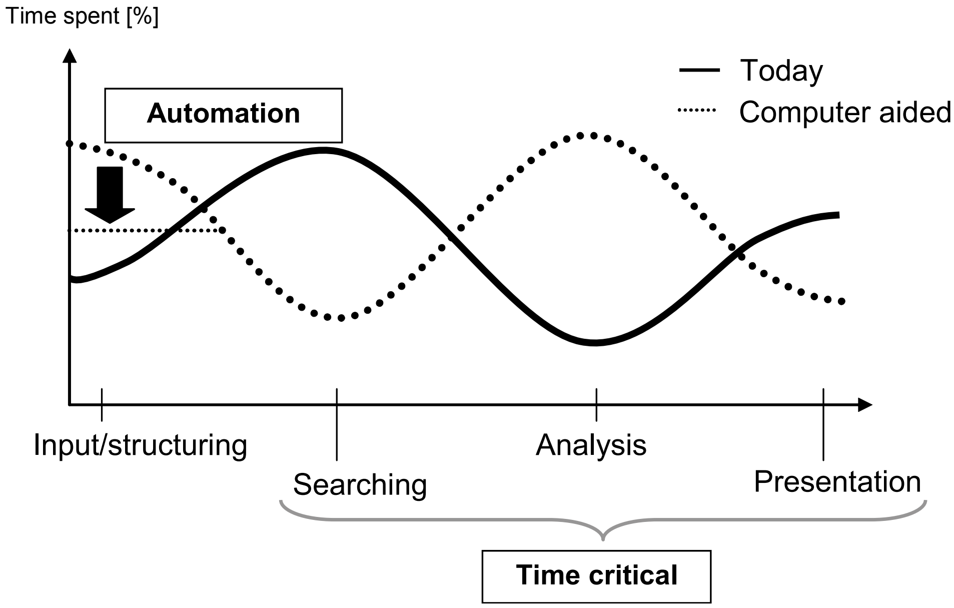

Figure 1: Theoretical intelligence workflow. The solid line shows the amount of time spent in the search, analysis and presentation phases in a typical case of today, whereas the dotted line shows how much time may be spent in each phase if semantic techniques could be used. The planning phase is not included in the figure, but is expected to take about the same amount of time using semantic techniques as it is today.

using a tool such as the one described in Section 6.4. It is, however, likely that it will still be necessary to have humans participating in this process: language technology is not yet mature enough to be able to produce semantic tags with enough certainty. Since the consequences of making wrong decisions based on the processing of the semantic tags could be severe, it is doubtful whether complete automation of this process would even be desirable.

Another benefit is that if information is structured, it is easier to see how new information would affect a given situation. If structured information is added, the new information might cause a chain of reactions that may lead to new statements that will drastically change the situation picture. For example, a list of potential threats that is dynamically updated when new input is added may be very useful in time-critical situations.

To put our work into perspective, it should be contrasted and compared to the research and development efforts currently undertaken in support of the Swedish military intelligence function [26]. This work is done in close cooperation with intelligence personnel who have been continuously participating using action research methods [9]. That is, analysts and researchers have been working collaboratively in order to improve intelligence work procedures through continuous reflection on, and adjustment of, the actions taking place in the actual intelligence unit (in this case, the basic data used for improvement primarily consisted of interviews, oral reflection, diaries from field studies, and document analyses). Also, it should be emphasized that these undertakings have been published openly and, hence, contrasted to work performed throughout the wider intelligence community. In summary, we believe that the standards set up by the described research project are 1) based on highly topical and relevant judgments regarding the Swedish intelligence function, and 2) being compared to and up-to-date with similar undertakings in the rest of the world.

In even more recent work, the same authors elaborate on the multifaceted and somewhat vague concept of an “intelligence architecture” [27]. In this work, the authors’ long-term goal is to integrate the human and the technological aspects of intelligence work. Although resisting formal analysis and definition, the ISTAR (intelligence, surveillance, target acquisition, reconnaissance) concept provides precisely this, i.e., a possibility to overcome the problem of uniting widely differing pieces of information by careful integration of human work and technical compositions. A conclusion made is that there is little need for automated quantitative processes in the intelligence domain, largely because knowledge production is politically and command informed rather than being the result of a formalized objective process. We support this viewpoint and agree upon that the human should always be in command when it comes to intelligence work, but notice that it is still the case that intelligence work will be, and need to be, computer-assisted. That is, in order to perform their day-to-day work analysts use, e.g., ordinary spreadsheets, word processors, tools used for dissemination of results, as well as more sophisticated tools such as the Analyst's Workstation[2] suite that can be used to, e.g., view and enlighten, possibly large, pieces of intelligence information. Given these two perspectives, we take the view that intelligence technology ought to be looked upon and constructed in the form of tools, i.e., a variety of supporting tools that the analyst uses to facilitate his/her day-to-day work. Moreover, we believe that the development of this "analyst toolbox" should be thought of as a continuously ongoing process involving practitioners, researchers and engineers. On the one hand, tool development needs to be informed by actual intelligence work procedures. On the other hand, intelligence work needs to be informed by current information technology trends to avoid getting stuck with old technology. In the following sections, we will describe some preliminary thoughts on a number of partly novel tools that could be part of such an "analyst toolbox."

5 A Blackboard Architecture for Intelligence Analysis

The computer tools used by intelligence analysts need to be connected together in what is called an intelligence architecture.

A blackboard architecture has previously been suggested to be suitable for collaborative intelligence analysis work [36]. We agree with this and further believe that the blackboard architecture would be beneficial as the basis of our analyst-centric toolbox also in the case of a single analyst. A blackboard system serves as a shared memory space for a number of agents independently contributing to the solution of a common problem. As depicted in Figure 2, the system consists of a blackboard, knowledge sources (the agents) and a control mechanism:

- The blackboard is the global "memory wall" where all information relevant for the problem and its solution is published. This normally includes the problem formulation itself, partial solutions, suggestions and data. In the case of our analyst-centric toolbox the blackboard would comprise the outputs of each tool.

- The knowledge sources are the components that publish new information to the blackboard. Taking the state of the blackboard as input they decide whether or not they can contribute. The tools in our toolbox can be seen as different knowledge sources publishing information on behalf of one or many analysts.

- The control mechanism is responsible for making sure that the different knowledge sources collaborate optimally. This task includes guiding the workflow, solving publishing conflicts and suggesting collaborative constellations in order to solve complex problems. As for our intelligence toolbox, it is apparent that in the case of many analysts working with different tools simultaneously there is a need for a controlling mechanism to optimize workflow and avoid potential conflicts. However, even in the

[2] See http://www.i2.co.uk/ for product information.

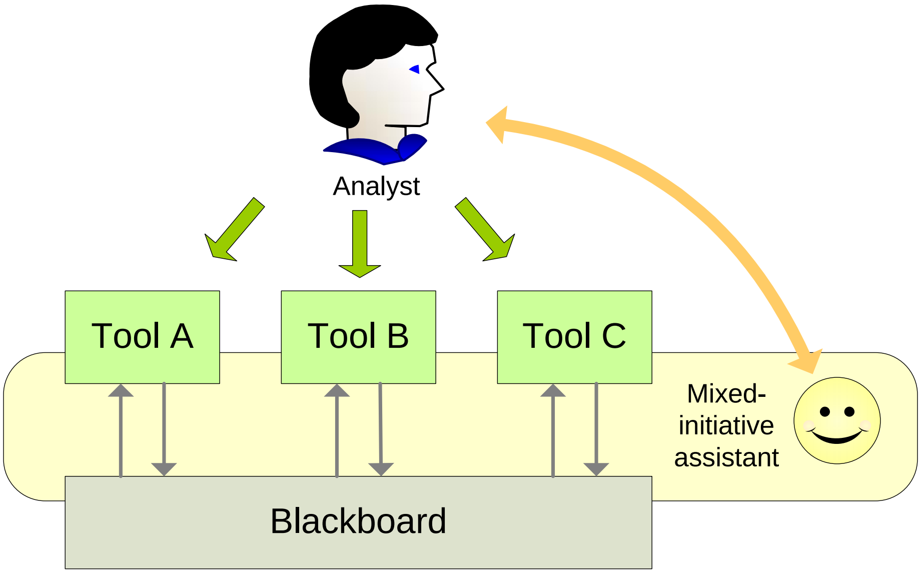

Figure 2: Schematic visualization of a blackboard architecture with a mixed-initiative assistant that helps the user determine what tool to use.

case of a single analyst, the complexity of today's information management calls for an assisting control agent guiding the analyst to areas of the problem where he/she can make contributions and suggest which tool to use.

One way of constructing such an assistant is to use techniques from mixed-initiative interaction [16, 41]. Mixed-initiative interaction attempts to construct artificial intelligence tools that interact with the user in a way that is similar to the way that two humans interact. The emphasis is on dialogue and interaction between the user and the machine. The field attempts to make use of the best capabilities of both the computer, e.g., the capacity to perform vast amounts of calculations fast, and the human, e.g., the ability to reason and human intuition. By allowing the human to interrupt the computer and the computer to provide hints to the human, better reasoning is enabled.

Proposing the use of a mixed-initiative assistant for the intelligence toolbox raises questions on automation. As stated in the previous section, we believe that intelligence work in general lends itself poorly to automated processes. This is a rather imprecise statement as there are many different levels of automation. In [33], automation is divided into 10 levels, from completely manual processes (level 1) to full automation (level 10), see Table 1.

A mixed-initiative assistant can in theory act on the whole automation scale. In the context of our toolbox, we see a number of different tasks for the assistant with varying degrees of automation. The assistant can offer the analyst hints on which tool to use or operation to perform next. This would typically correspond to automation on level 3 or 4. The assistant could also perform operations on its own in the background, presenting relevant findings to the analyst when suitable. This would correspond to automation levels 5 to 10. However, the results of such actions should be approved by the analyst before written to the blackboard, alternatively be marked so that consumers of the information know it has automatic origin. An important factor when determining the degree of automation is the level of operation. Due to sharper time constraints, a tactical intelligence product is more likely to contain traces of higher automation levels than a strategic one.

Automation level	Automation description
1	The computer offers no assistance: human must take all decisions and actions.
2	The computer offers a complete set of decisions/action alternatives, or
3	narrows the selection down to a few, or
4	suggests one alternative, and
5	executes that suggestion if the human approves, or
6	allows the human restricted time to veto before automatic execution, or
7	executes automatically, then necessarily informs humans, and
8	informs the human only if asked, or
9	informs the human only if it, the computer decides to.
10	The computer decides everything and acts autonomously, ignoring the human.

Table 1: Levels of automation according to [33].

6 Analysis Tools

In this section, we briefly describe several components of our proposed intelligence analysis toolset. The selection of tools presented is preliminary and represent our current work. We envision that our intelligence toolbox will contain a large number of other tools as well.

In order to be use-worthy, computer tools for intelligence analysis must both have good, intuitive user interfaces and provide functionality that is actually useful. One way of constructing such tools is to use techniques from mixed-initiative interaction, as described in Section 5.

6.1 Report Organization and Visualization

A useful combination of information search and information analysis is to sort incoming reports and sensor observations and present different views of them to the user. By grouping similar reports together and attaching a meaningful label to the groups, the user's situational awareness is increased. Through the years FOI has developed several different methods for doing this for different applications. Example areas worked on include clustering of sensor observations of tanks [2], labeling of observed groups according to what capability they have to attack protection-objects [38], and constructing a view of reports based on what future enemy course of action they might indicate [39]. Figure 3 shows an example of the last application, where reports are visualized on the right (the map and the textual description) and possible future events are shown on the left. After the user has clicked on an event, the system searches for those reports that indicate that the event is occurring and displays them to the user. In this way, instead of having to read all reports, the user can quickly scan only those reports that are relevant to the event that they are currently considering, which leads to faster situational awareness.

Another approach is to determine what resources the opponents have and at what locations they are placed. By doing this, we can determine what capabilities they have to harm us or civilian targets, and plan our actions accordingly [37].

6.2 Impactorium

Impactorium, see Figure 4, is a tool helping users estimate the probabilities of various events that might occur in the future and which will have an impact on the user. It is based on the Impact matrix concept, where events of interest are visualized in a 2 × 2 matrix where the axes represent a priori probability and impact, respectively. In [39], we described a dynamic Impact matrix that has a semi-automatic coupling to real-world

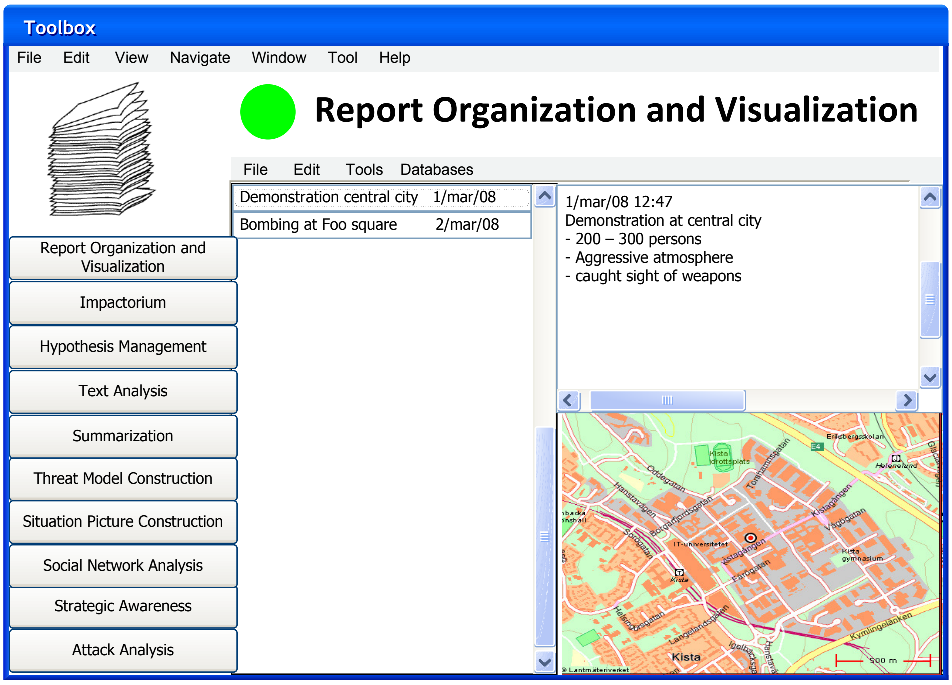
Toolbox
File Edit View Navigate Window Tool Help
Report Organization and Visualization
File Edit Tools Databases
Report Organization and Visualization
Impactorium
Hypothesis Management
Text Analysis
Summarization
Threat Model Construction
Situation Picture Construction
Social Network Analysis
Strategic Awareness
Attack Analysis
Kista
500 m

Demonstration central city	1/mar/08
Bombing at Foo square	2/mar/08

1/mar/08 12:47
Demonstration at central city
- 200 – 300 persons
- Aggressive atmosphere
- caught sight of weapons

Figure 3: User interface of a tool for sorting and clustering reports.

events that change the probability that an event is occurring or will occur in the near future. This is done by coupling incoming HUMINT reports to indicators that give a high-level description of what they are about. The indicators are assumed to be associated to the reports either automatically or, most likely, by human operators that monitor sensors and incoming reports.

Indicators are generated using a Bayesian network linked to hypotheses about the realization of events; that is, the indicators, together, give different probabilities that a certain event will happen. When the probability of an event changes, it is indicated in the matrix by changing the color and size of the circle displayed next to the event. Thus, by looking at the matrix, it is possible for the user to immediately spot events that are about to occur, hopefully leading to increased situational awareness and giving the user the opportunity to act proactively.

Currently, Impactorium only support model-based threat analysis. This means that the indicators as well as the Bayesian network used to calculate the event probability based on the values of the indicators must be specified by a subject-matter expert beforehand. Another approach to generating threat models for use in Impactorium would be to use statistical methods for learning accurate threat models from a large set of data. This, however, requires large amounts of data, which is not available for our problem domain of interest. There are two drawbacks to using the model-based approach: specifying the models requires a large amount of human effort, and it is not possible to automatically discover new threats for which there are no models. We are tackling these difficulties in several ways. As outlined in Section 6.5, we will attempt to construct tools that help the users construct appropriate threat models. We are also working on a method for automatically updating the Bayesian network used in a threat model if the user determines that the value calculated for the probability is wrong. Future versions of Impactorium will also include functionality for clustering observation reports that could be used to semi-automatically learn new threat models.

We have done user experiments on the Impactorium tool both on tactical and strategic

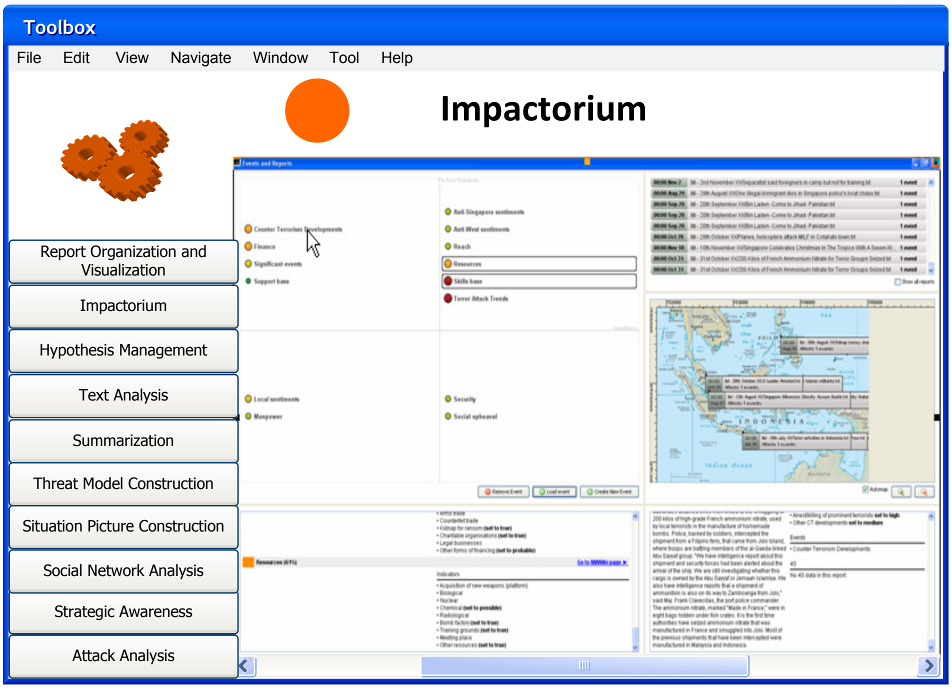

Figure 4: User interface of the Impactorium tool.

levels [11, 25]. These experiments indicate that the tool is useful and helps the decision-makers achieve better situational awareness. When showing the prototype to users who have been in service, we have received several interesting suggestions for other uses of the tool. The impact matrix could be used to display information about what NGO's and other civilian actors in the area are doing. It could also be useful for training and for debriefing, where it would be useful to replay the events that have taken place and see what other actions could have been taken. Interestingly, a company commander also suggested that he would have liked to have had access while in the field to an annotated picture of the Impactorium display, where analysts at the FHQ (force headquarters) had added extra information.

In one experiment that we performed we gave some people the Impact matrix tool while others were provided with a program displaying all the reports and allowing the user to view them (essentially, this was the right-hand side of the full Impactorium interface). In this setting, the user who only got the reports was able to spot one upcoming event, but missed that there were also indicators present for one more high-probability event and two medium probability events. This was an indication to us saying that the Impact matrix program fulfils an important role.

In addition to being used by Impactorium and other tools, the indicators mentioned above are also useful by themselves. By constructing networks that connects those indicators that are associated to the same reports, or those reports that are connected to the same indicator, the user gets yet another way of navigating through the report database. By connecting these networks to people, places, and organizations, it is also possible to use social network analysis (SNA) methods to gain more information about the most important indicators and reports.

Social network analysis [43] and visualization is an essential capability for the battle groups. By analyzing the reports received, it is possible to infer the communication patterns of the actors in the area of responsibility. Community detection algorithms could be used to discover new groupings among them. Robust Bayesian methods could be used to handle uncertain network data. See Section 6.7 for a more in-depth outline of the SNA tool that we have in mind for intelligence analysis.

6.3 Hypothesis Management

Situational awareness does not mean that the user has a clear picture of what is happening now and in the near future. Instead, in all but the simplest situations, the user will have several different hypotheses regarding what is happening and what will happen next. The reason for this is the uncertainties that are inherent in the information fusion process. The users thus need to be aware of how to handle uncertainty, and also need tools that allow them to visualize and reason about their hypotheses. Since the human often can only mentally keep track of a small number of differences between the hypotheses, there is a need for tools that remind the user of the most important differences between the hypotheses, and updates them according to the new information that is constantly fed into the system both by the information gathering resources and by the analysis tools. A mixed-initiative agent that points out the most important details to the user should be implemented in the decision support system. Capability similar to the one presented in [34] for determining when the hypotheses about the situation picture has changed sufficiently and replanning needs to be done, should also be included.

6.4 Semantic Tagging by Knowledge Extraction from Data

The semantic web is the term used for the vision of making Internet content interpretable by machines. The semantic web helps computers gain better understanding of what information really means. Subsequently, when the information is understandable by computers, the computers would also be able to infer new facts. Hence, the semantic web concept, where information can be interpreted by machines, may not only improve the Internet; it may also be useful in other information sharing domains such as information sharing within the armed forces.

In order to make full use of the semantic techniques, it is necessary for the content of KB's to be semantically tagged. Although there is still much work to be done, research in this area is active and making good progress [42], and performing semantic tagging will inevitably be a major part of intelligence analysts' work in the future. Hence, automating some parts of this process would be very useful. One approach to doing so is to use text analysis techniques to perform entity extraction and present a list of entities found in the document to the user who is doing the semantic annotation. A more challenging problem is to also find relations in the document, and to add semantically marked links to the intelligence document. Figures 5 and 6 show a simple prototype system.

In addition to the annotation of each document, it is also important for the intelligence analysis assistant program to be able to analyze several documents at the same time. For this, summarization techniques from natural language processing could be used, see, e.g., [15] for a generic approach. It should also be possible to obtain lists of the extracted entities that are present in the different documents and fuse these to provide a "situation picture" of what objects are referred to in the document collection. In addition to standard techniques for this, we are also investigating the use of topic models [35] to fuse text documents with sensor data and semantically annotating the result.

As mentioned in Section 4, we do not believe it will be possible (or desirable) to automate the semantic tagging process. However, computer tools acting as assistants to the human, suggesting tags for documents, would provide a valuable enhancement of the intelligence input/structuring process.

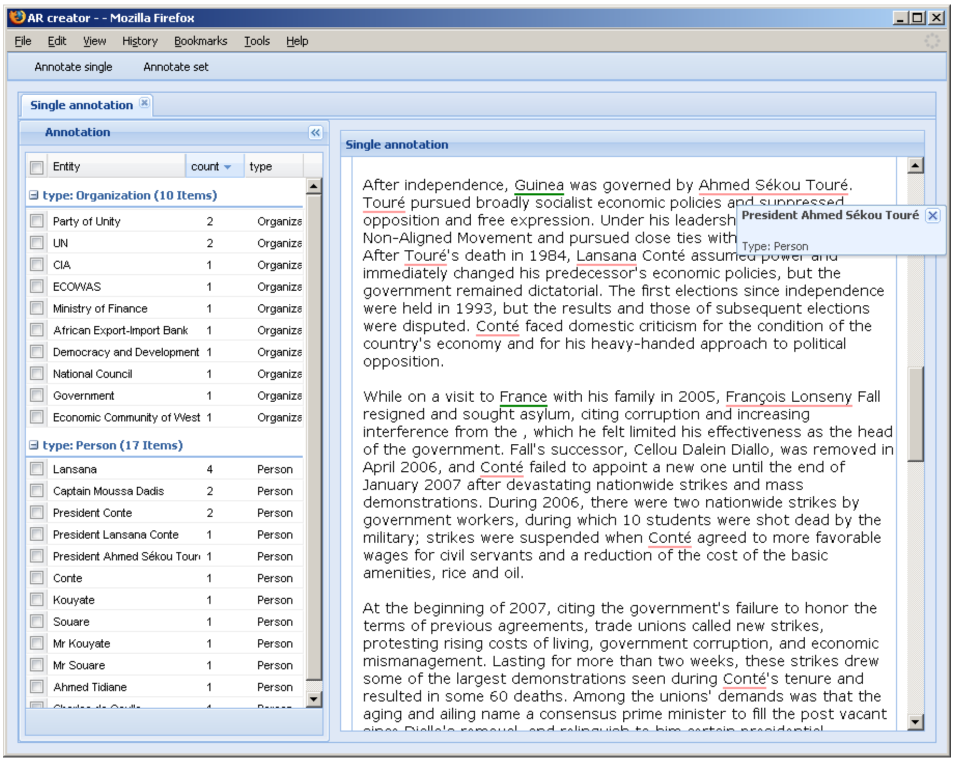

□	Entity	count	type
	type: Organization (10 Items)		
□	Party of Unity	2	Organiza
□	UN	2	Organiza
□	CIA	1	Organiza
□	ECOWAS	1	Organiza
□	Ministry of Finance	1	Organiza
□	African Export-Import Bank	1	Organiza
□	Democracy and Development	1	Organiza
□	National Council	1	Organiza
□	Government	1	Organiza
□	Economic Community of West	1	Organiza
	type: Person (17 Items)		
□	Lansana	4	Person
□	Captain Moussa Dadis	2	Person
□	President Conte	2	Person
□	President Lansana Conte	1	Person
□	President Ahmed Sékou Tour	1	Person
□	Conte	1	Person
□	Kouyate	1	Person
□	Souare	1	Person
□	Mr Kouyate	1	Person
□	Mr Souare	1	Person
□	Ahmed Tidiane	1	Person

After independence, Guinea was governed by Ahmed Sékou Touré. Touré pursued broadly socialist economic policies and suppressed opposition and free expression. Under his leadersh Non-Aligned Movement and pursued close ties with After Touré's death in 1984, Lansana Conté assumed power and immediately changed his predecessor's economic policies, but the government remained dictatorial. The first elections since independence were held in 1993, but the results and those of subsequent elections were disputed. Conté faced domestic criticism for the condition of the country's economy and for his heavy-handed approach to political opposition.

President Ahmed Sékou Touré

Type: Person

While on a visit to France with his family in 2005, François Lonseny Fall resigned and sought asylum, citing corruption and increasing interference from the , which he felt limited his effectiveness as the head of the government. Fall's successor, Cellou Dalein Diallo, was removed in April 2006, and Conté failed to appoint a new one until the end of January 2007 after devastating nationwide strikes and mass demonstrations. During 2006, there were two nationwide strikes by government workers, during which 10 students were shot dead by the military; strikes were suspended when Conté agreed to more favorable wages for civil servants and a reduction of the cost of the basic amenities, rice and oil.

At the beginning of 2007, citing the government's failure to honor the terms of previous agreements, trade unions called new strikes, protesting rising costs of living, government corruption, and economic mismanagement. Lasting for more than two weeks, these strikes drew some of the largest demonstrations seen during Conté's tenure and resulted in some 60 deaths. Among the unions' demands was that the aging and ailing name a consensus prime minister to fill the post vacant

Figure 5: A document where entity extraction has been made. Entities recognized by the used ontology have been marked with colors depending on their meaning, e.g., persons have been underlined with red and organizations have been underlined with green.

6.5 Threat Model Construction Assistant

As outlined in Section 4, fusion tools need models in order to work. For example, the Impactorium tool described in Section 6.2 needs Bayesian networks that describe how indicators are used to estimate the probability of various events of interest. Constructing these models is an important, and difficult, part of intelligence analysts' work. We think that it is possible to use case-based reasoning to help the humans in this process. Case-based reasoning relies on a case database consisting of previously seen solutions to problems. For a recent survey of case-based reasoning and, in particular, its relationship to analogical human reasoning, see [22].

In our application, a case would correspond to a "situation," as described, for instance, by the set of active indicators, see Figure 7. A solution would correspond to a particular model to use in the fusion tools, for instance, the Bayesian networks to be included in the Impactorium tool. The current situation is compared to those in the case database, and the best matches are extracted and presented to the user who can choose which models to include. In order to make the match, situations must be compared with each other and a similarity measure calculated. For attributes adopting numerical values, it is easy to calculate the similarity distance. For other kinds of attributes, other models need to be used.

To determine the extent to which this process can be automated, it is necessary to first make an adequate description of a situation. For this, it will be necessary to define an ontology or information model that is rich enough to be able to distinguish between different situations. Some first steps towards this have been taken in [17]. Construction of a relevant similarity measure between situations is also a challenge. A problem for this approach is of

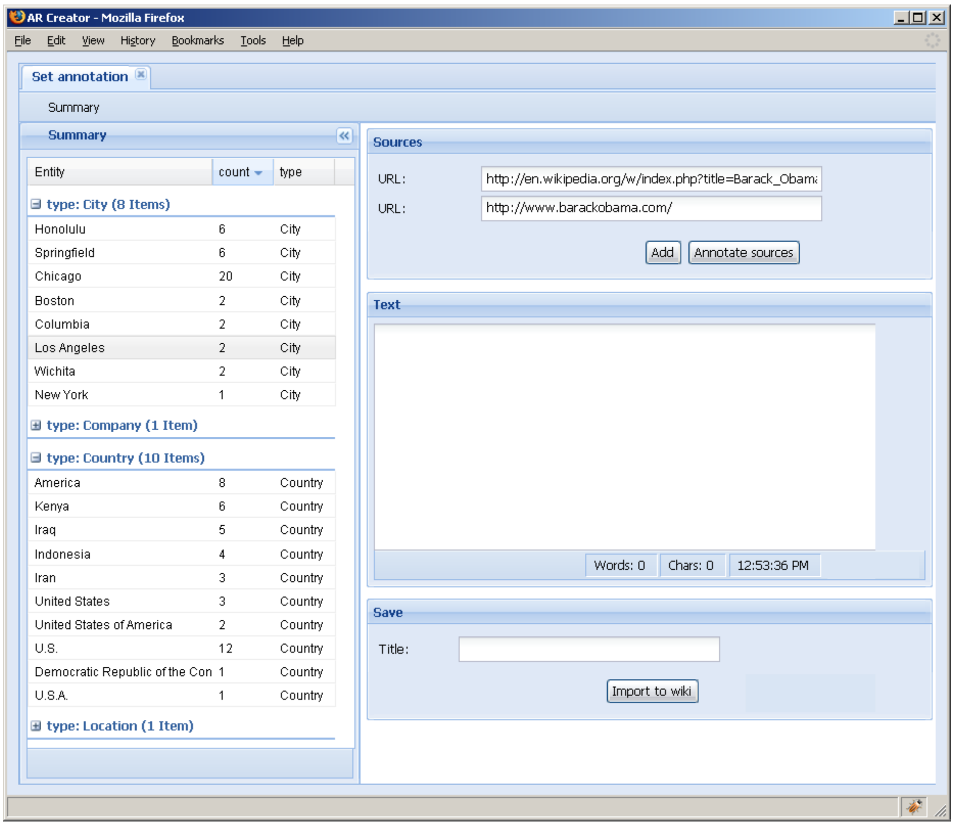

Entity	count	type
type: City (8 Items)		
Honolulu	6	City
Springfield	6	City
Chicago	20	City
Boston	2	City
Columbia	2	City
Los Angeles	2	City
Wichita	2	City
New York	1	City
type: Company (1 Item)		
type: Country (10 Items)		
America	8	Country
Kenya	6	Country
Iraq	5	Country
Indonesia	4	Country
Iran	3	Country
United States	3	Country
United States of America	2	Country
U.S.	12	Country
Democratic Republic of the Con	1	Country
U.S.A.	1	Country
type: Location (1 Item)		

URL: http://en.wikipedia.org/w/index.php?title=Barack_Obama

URL: http://www.barackobama.com/

Words: 0 Chars: 0 12:53:36 PM

Title:

Figure 6: The figure shows how an interface to be used by an intelligence analyst looking at a number of intelligence documents could look like. The panel on the left shows aggregate statistics about entities found in documents—information that could be of use when writing an intelligence report.

course that the number of cases in the case data base will initially be low. Situations from previous missions should be reused, so that we can make use of, for example, lessons learned in Kosovo for the mission in Afghanistan. Developing a case-based reasoning mechanism that is sophisticated enough to do this is a major research challenge.

6.6 Situation Picture Construction Assistant

Network-based defense decision-makers and analysts are given the opportunity to make use of a wide range of information services. For a specific mission in a specific situation, however, it is only a small part of these services that are useful. A choice needs to be made regarding what services to use. Since each service will provide a "building-block" of a situation picture, it is possible to create different views of the situation by making use of different services. In order to support an officer in the selection of information, one could make use of techniques from case-based reasoning. Good solutions for what set of services to use are stored in a case database along with a description of the overall situation when the set of services were selected. Factors such as the role of the decision-maker and the status of own operations should also be included in the case description, as well as external factors such as mood and cultural status in the region.

User-tailored situation-views could also be constructed using other methods. An example might look something like this. An intelligence analyst is searching for information to

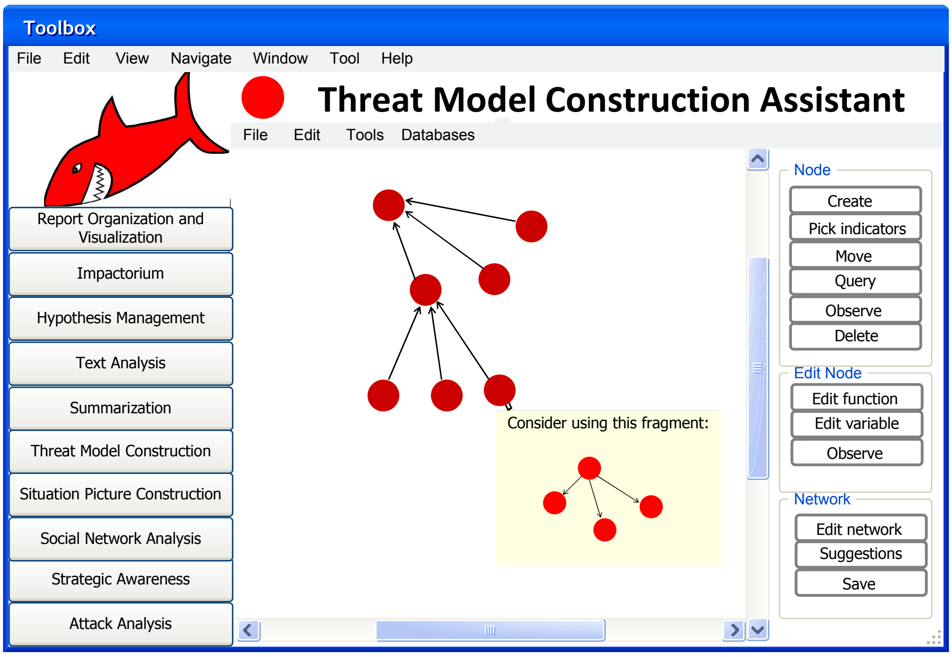

Figure 7: Visualization of a particular situation obtained from the case database using the model construction assistant.

respond to an RFI (request for information). In a given moment, the analyst is looking at the information provided by the KB about the warlord X. The system should then show related information that it believes the user could be interested in, at the side of the computer screen. For instance, if the context within which the analyst is working relates to smuggling, information about boats owned by X or an associate, and that have been seen near a border could be shown. A more sophisticated example would be to show information about a boat moving anomalously and which is linked to a subordinate of X. The idea is similar to the recommendation system used by many online stores, see, e.g., [1, 19], but would require more advanced methods.

Another example is if the analyst is interested in criminal activities in the theatre of operations. The side-screen will then display an aggregated statistical view of the entire area showing crime statistics. The user can choose to zoom in on parts of the data by, e.g., looking at a map where crimes committed by a member of a certain ethnical group having a member of another ethnical group as victim have been committed.

6.7 Social Network Visualization and Analysis

Visualization and analysis of network data is an essential part of intelligence analysis. By forming networks of interesting people, organizations, places and events connected with links that indicate connection, we can create conceptual images that facilitate understanding. Network analysis is also important in mission debriefing, or after an exercise: by studying how the information has moved in the organization we can find opportunities to improve management. The field of social network analysis (SNA) provides tools and methods for such analysis.

SNA is a set of powerful techniques, to identify social roles, important groups and hidden organizational structures. Correlation of observed data about individuals, things, places, memberships, etc., may be used to detect organized crime or terrorist cells and networks through the observation of hidden relations and co-occurrences. This methodology assumes that the ways the members of a group can and do communicate with each other affect some important properties of that group. Figuring out nested business connections across a set of known individuals or organizations is one application of SNA. Since not all people who have had contacts with a criminal are criminals themselves, there is a need for techniques which can filter out those whose contacts with known or suspected criminal individuals are either frequent or match stored patterns of suspicious behavior. One issue thus deals with how one can automatically estimate which people among a very large community, that have been "transitively" in contact with each other, need to be investigated further and who do not.

In previous work [12] we have implemented a tool for distributed, collaborative network analysis and visualization. The user interface of this tool is shown in Figure 8. The tool allows users situated in different geographical locations to collaborate on the same network data. The toolset also includes components for filtering large amount of data for relevant network data, and for merging network data from several sources into a single network to analyze.

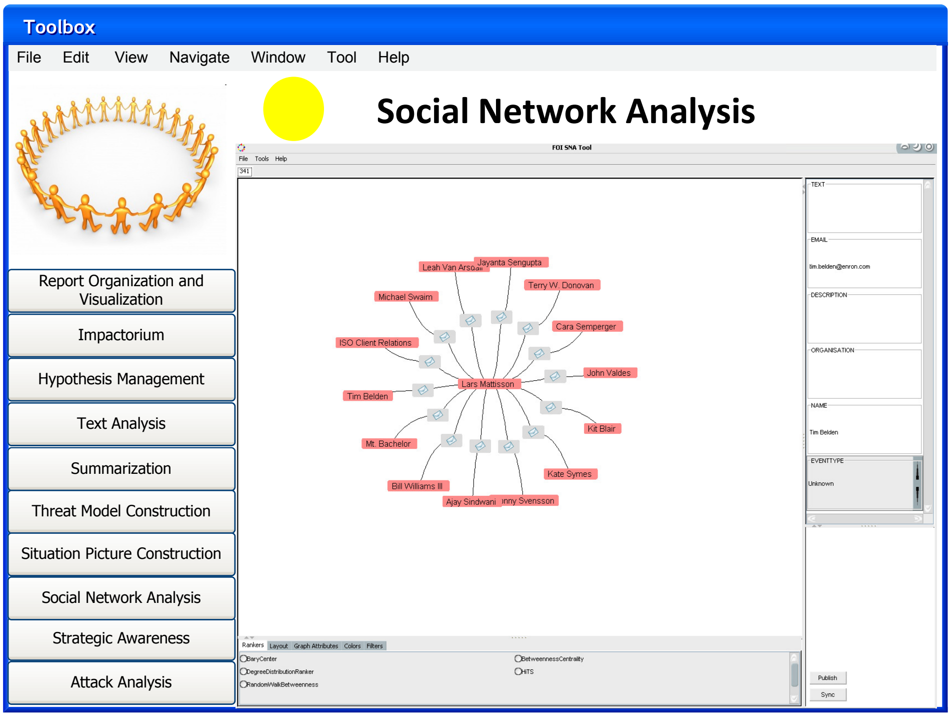

EMAIL
tim.belden@enron.com

NAME
Tim Belden

EVENTTYPE
Unknown

Figure 8: User interface of a tool for distributed, collaborative social network visualization and analysis.

6.8 Attack Analyzer

Visualizing and modeling possible attacks towards a goal is a difficult but important task. If such attacks can be modeled in a suitable way the user can gain a deeper understanding of attacks which makes it possible to design and evaluate countermeasures to prevent some of the attacks. In this tool we use a formalism called attack trees [32] to model possible attacks. Attack trees provide a formal, methodical way of describing and modeling possible threats based on various attacks.

In the "attack analyzer," depicted in Figure 10, each attack is graphically represented in a tree structure with the goal of the attack as the root node. Children of a node are refinements of this goal and leafs represent attacks that can no longer be refined.

There are two types of nodes in an attack tree, AND-nodes and OR-nodes. Satisfying a node means either satisfying all predecessor nodes (AND-node) or satisfying some predecessor node (OR-node). When the root is satisfied, the attack represented by the root is complete.

Figure 9 shows an example of a threat modeled as an attack tree. In the figure, all nodes that are not marked as AND-nodes are OR-nodes. The goal of the threat is to kill a person named X. An attacker can kill X by strangling him, stabbing him, shooting him, poisoning him or bombing him. To be able to stab X, the attacker needs a knife and he/she needs to get close to X. To get close enough to X, assuming that X is someone who is under high surveillance, an attacker has to either work in X's staff or he/she needs to work at a location that X visits. The only possible locations are the hairdresser, X's favorite restaurant or the gym where X works out.

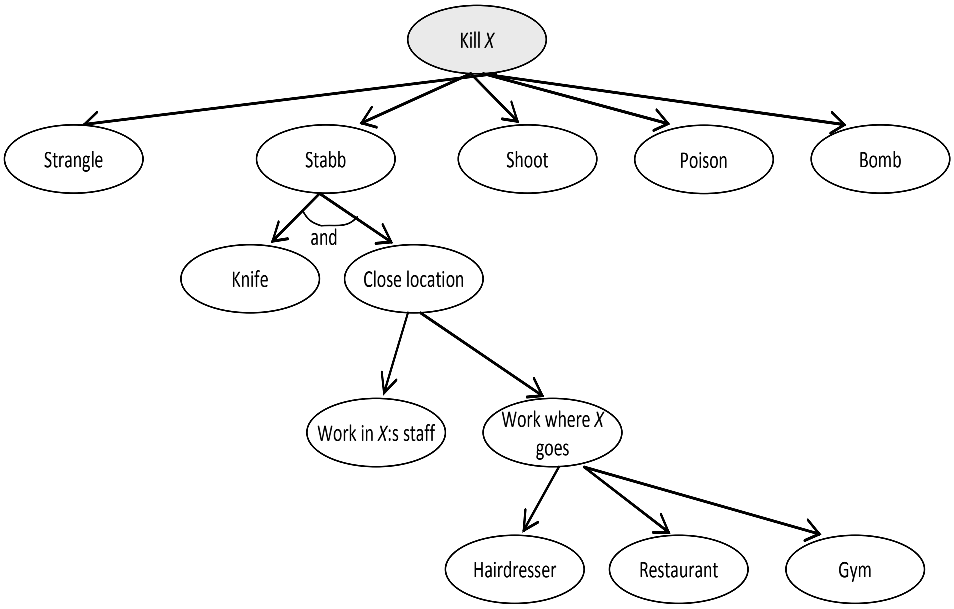

Figure 9: An (incomplete) attack tree where the goal is to kill X.

Once the attack is modeled as an attack tree, the tree can be analyzed. The analysis is performed in two steps: first all paths leading to a root are extracted from the model and then each path is analyzed. A path consists of a set of attack components that are needed to perform the actual attack. Each attack component can be assigned an attribute. An attribute could for example be the possibility/impossibility of accomplishing an attack component. In Figure 9 we could argue that it is possible that someone who wants to stab X works at his favorite restaurant or at the gym but since his hairdresser is his uncle we can assign the value "impossible" to the attack component representing X's hairdresser.

Another example of attributes is whether or not special tools are required to accomplish an attack component. Depending on the attribute, the analysis can answer a number of different questions about the attack tree, such as "What is the minimum cost of performing an attack?" and "What is the minimum required skill level for an attack?"

For the paths in the attack tree, possible countermeasures can be evaluated and designed. For example, in Figure 9 a countermeasure to the attack where X get killed by a gunshot is that X can wear a bulletproof vest.

The formal representation of an attack tree enables tools to both create and analyze possible threats. A good tool support is required since the attack trees can become large and complex: a full attack tree may contain thousands of different paths all leading to completion of the attack.

The "attack analyzer" stores all previous models of attacks. Whenever a new attack tree is initiated the tool indicates whether there exist previously constructed trees that the user can view. This is useful to assure that information about previous attacks is included in newly constructed attacks as well.

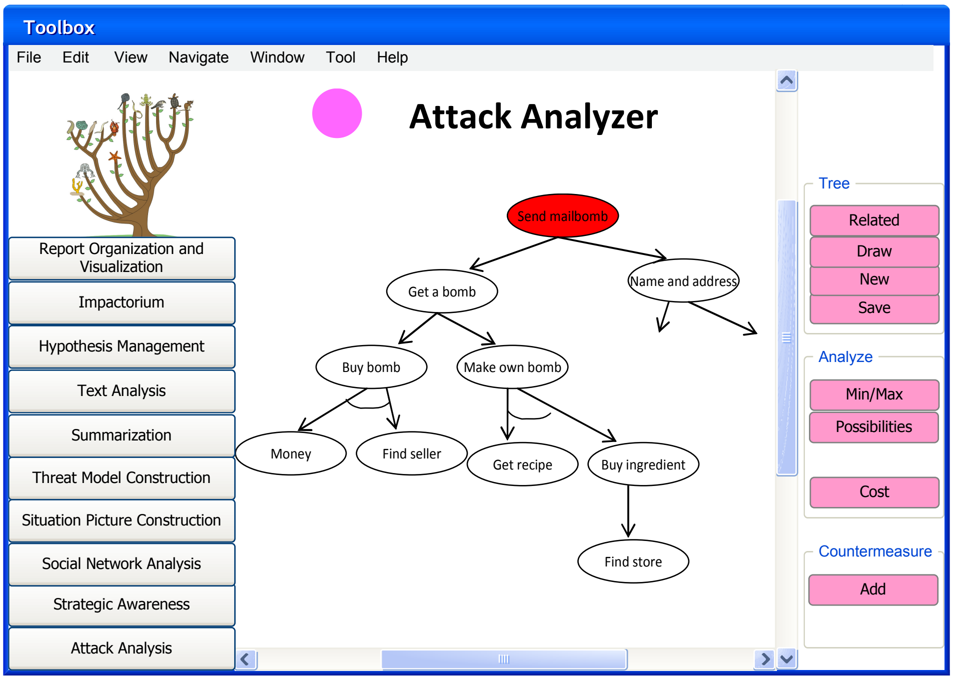

Figure 10: The "attack analyzer" makes it possible for the analyst to model and understand the severity of possible attacks.

6.9 Strategic Awareness Assistant

Military intelligence work ought to be performed within a "system of systems" consisting of several integrated intelligence processes. Each process functions on its own while it is at the same time interacting with other intelligence processes at various levels. These disparate, but hopefully well-integrated, processes are carried out within different phases of the intelligence cycle using different timescales and can be coarsely divided into three levels: tactical, operational, and strategic. The processes utilize results from each-other in order to provide decision-makers with timely and high-quality basic data for decision-making on all levels based on both short-term tactics and long-term general assumptions [40].

The "strategic awareness assistant" helps analysts' working at the strategic level by giving "rational advice" based on actors, actions, utility theory and the game-theoretic equilibrium concept [23]. At the strategic level, the intellectual part of the analyst work is directed towards identifying potential actors of interest, their possible courses of action and the probabilities associated with these courses of action. According to intelligence personnel that we have been communicating with, one usually tries to assess a limited number of actors and actions in order to come up with a small number of potential scenarios to reason further about.

The frontend of our application involves a graphical user interface, see Figure 11, letting the analyst list actors, potential courses of action, possible outcomes, and their associated probabilities. It should be noted that these preparatory steps that are routinely taken by intelligence personnel conform closely to traditional decision theory, see, e.g., [29], which, hence, makes way for the development of algorithms suitable for a computer-based decision support tool. The backbone of our application lets the analyst perform probability calculations and visualize the strategic awareness obtained in the form of mixed strategy Nash equilibria. To make it a useful tool in practice, however, would require extensive user experiments and creative solutions when it comes to visualization.

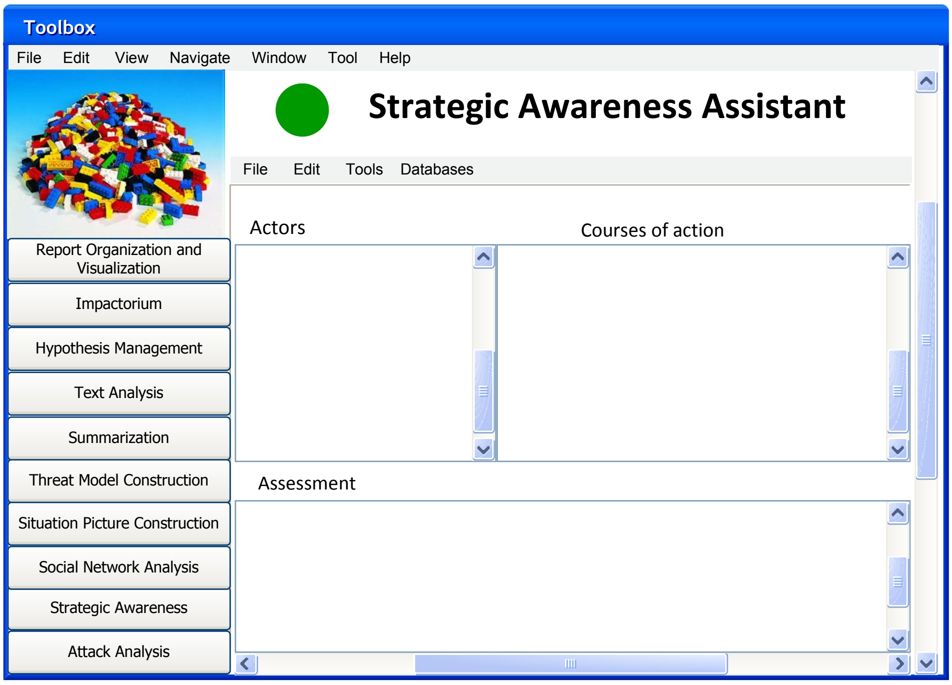

Figure 11: The strategic awareness assistant lets the analyst make game-theoretic assessments based on interesting actors and these actors' potential actions.

On tactical level, our tool can be thought of as an ordinary decision-theoretic tool, a game against nature where the sum of the utilities provides a non-controversial solution [20]. Here, there is no time to let one's own thinking depend on recursive reasoning about the opponent's mindset. Rather, one act based on skill, doctrine, or trained procedure. Probably one wishes to reason about a few possible predetermined worlds where ordinary probabilistic methods can be used to efficiently calculate optimal solutions given situation-specific parameters. We do not think that our tool is suitable at the tactical level due to the available time and the level of complexity, but think that the decision-theoretic game against nature serves as a good starting point for our explanation. At the strategic level, on the other hand, reasoning about reasoning must be thought of as the foremost intellectual ingredient of the analysts' work. Here, intelligence work turns into a game resembling that of poker, i.e., a game containing uncertainty caused both by nature and by willful-thinking opponent players. Since all players must be assumed to be able to make the same calculations we cannot solely optimize and calculate some sort of "best solution" since we would easily be outperformed if we did so. Instead, the Nash equilibrium concept provides a non-trivial, but rationally sound, "strategic awareness" concept.

To obtain awareness in the form of, potentially several, mixed strategy equilibria requires a moment of thought. At first sight, it might feel uncomfortable to obtain the "optimal solution" in the form of probability distributions, but this should not be taken to be the case. Instead, it makes perfect sense to think of "awareness" in terms of equilibria: a probability measure gives you (rational) awareness rather than a precise answer and is precisely the kind of mathematical measure that we are looking for. Also, at the strategic level the difference between the application of force and the threat of actually using force must be made [31]. It is one thing to actually use force or other measures to conduct operations successfully—yet another to use the potential of such measures in order to enforce goals by deterrence. Hence, it makes perfect sense to reason about these kinds of awareness concepts. As said, however, it remains to make the equilibrium concept easy to grasp in order to develop a really useful computer tool.

A future enhancement of the “strategic awareness assistant” could be to use the notion of a Bayesian game [14] in order to allow for commanders to have different views of the underlying game model. For details regarding how this technique can be applied, see [6] and [7] for an example scenario and a more detailed description respectively.

7 Discussion

In all the tools discussed above, handling uncertain information is important. Information that comes from sensors, humans, or the web needs to be marked not only with relevant tags and indicators from an ontology, but also with the credibility of the reporting source and the estimated certainty of the information. These uncertainties further need to be taken into account by the analysis tools. For traditional information fusion tools, this is standard procedure, but for others the addition of uncertainty handling raises important new research questions. For example, how should uncertainty be handled in social network analysis? How should so-called “negative information” (that is, knowledge that something is not there, e.g., by having a sensor look for a car in an intersection and not seeing it) be handled? For a discussion on negative information and handling of uncertain threat models, see [21].

Visualization of uncertainties is also an important research question. For map-bound information, there exist standard procedures for doing this. In OOTW situations, however, it is the situation picture information that cannot be easily displayed on a map that is the most important. How should uncertainties in that information be presented to the user? An overview of some standard uncertainty visualization techniques can be found in [30].

There is also a need for inference engines that combine knowledge present in the KB with new, incoming reports and draw new conclusions. Such capability is of course needed in the tool that handles indicator management, but should also be a useful component in the hypotheses management tool.

It is vital for the force to have access to a good information system. The database should not only store results from analysis performed during the operation, but also background information about the area and relevant information from previous missions. Thus, the database should serve as an organizational memory, where commanders could look up information about situations that are similar to the current one and get advice about what to do. For this, it is likely that case-based reasoning technologies should prove useful.

References

[1] Gediminas Adomavicius and Alexander Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749, June 2005.

[2] Simon Ahlberg, Pontus Hörling, Katarina Johansson, Karsten Jöred, Hedvig Kjellström, Christian Mårtenson, Göran Neider, Johan Schubert, Pontus Svenson, Per Svensson, and Johan Walter. An information fusion demonstrator for tactical intelli-

gence processing in network-based defense. *Information Fusion*, 8(1):84–107, January 2007.

[3] Stefan Arnborg, Henrik Artman, Joel Brynielsson, and Klas Wallenius. Information awareness in command and control: Precision, quality, utility. In *Proceedings of the Third International Conference on Information Fusion (FUSION 2000)*, volume 2, pages ThB1/25–32, Paris, France, July 2000.

[4] Tomas Berg, Christian Mårtenson, and Pontus Svenson. Using text clustering for intelligence classification. In *Stockholm Contributions in Military-Technology 2007*, pages 23–34. National Defence College, Stockholm, Sweden, May 2008.

[5] James O. Berger. An overview of robust Bayesian analysis. *Test*, 3(1):5–124, June 1994.

[6] Joel Brynielsson. Using AI and games for decision support in command and control. *Decision Support Systems*, 43(4):1454–1463, August 2007.

[7] Joel Brynielsson and Stefan Arnborg. An information fusion game component. *Journal of Advances in Information Fusion*, 1(2):108–121, December 2006.

[8] Robert M. Clark. *Intelligence Analysis: A Target-Centric Approach*. CQ Press, Washington, District of Columbia, 2004.

[9] Louis Cohen, Lawrence Manion, and Keith Morrison. *Research Methods in Education*, chapter 14, pages 297–313. Routledge, London, United Kingdom, sixth edition, 2007.

[10] Thomas Dean and Mark Boddy. An analysis of time-dependent planning. In *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88)*, pages 49–54, Saint Paul, Minnesota, August 1988.

[11] Wong Rong Fah, Magdalene Selina Choo, and John Kho. Bayesian technique modelling tool for the risk assessment and horizon scanning domain? Manuscript, Defence Science and Technology Agency, Singapore, 2008.

[12] Luigi Ferrara, Christian Mårtenson, Pontus Svenson, Per Svensson, Justo Hidalgo, Anastasio Molano, and Anders L. Madsen. Integrating data sources and network analysis tools to support the fight against organized crime. In *Proceedings of the IEEE ISI 2008 International Workshops: PAISI, PACCF, and SOCO 2008*, pages 171–182, Taipei, Taiwan, June 2008.

[13] David L. Hall, James Llinas, Michael McNeese, and Tracy Mullen. A framework for dynamic hard/soft fusion. In *Proceedings of the 11th International Conference on Information Fusion (FUSION 2008)*, Cologne, Germany, June 30–July 3, 2008.

[14] John C. Harsanyi. Games with incomplete information played by “Bayesian” players. *Management Science*, 14(3,5,7):159–182, 320–334, 486–502, 1967–1968.

[15] Martin Hassel. *Resource Lean and Portable Automatic Text Summarization*. PhD thesis, School of Computer Science and Communication, Royal Institute of Technology, Stockholm, Sweden, June 2007.

[16] Marti A. Hearst, James F. Allen, Curry I. Guinn, and Eric Horvitz. Trends & controversies: Mixed-initiative interaction. *IEEE Intelligent Systems*, 14(5):14–23, September/October 1999.

[17] Pontus Hörling, Mikael Lundin, Christian Mårtenson, and Pontus Svenson. Informationsmodeller och indikatorer för situations- och hotbeskrivning [Information models and indicators for situation and threat description]. Technical Report FOI-R--2535-SE, Swedish Defence Research Agency, June 2008.

[18] Martin E. Liggins, David L. Hall, and James Llinas, editors. *Handbook of Multisensor Data Fusion: Theory and Practice*. Electrical Engineering and Applied Signal Processing Series. CRC Press, Boca Raton, Florida, second edition, 2009.

[19] Greg Linden, Brent Smith, and Jeremy York. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1):76–80, January/February 2003.

[20] R. Duncan Luce and Howard Raiffa. *Games and Decisions*. John Wiley & Sons, New York, 1957.

[21] Mikael Lundin, Pontus Hörling, and Pontus Svenson. Uncertainty modelling for threat analysis. In *Proceedings of the 14th International Command and Control Research and Technology Symposium (14th ICCRTS)*, Washington, District of Columbia, June 2009.

[22] Ramon López de Mántaras, David McSherry, Derek Bridge, David Leake, Barry Smyth, Susan Craw, Boi Faltings, Mary Lou Maher, Michael T. Cox, Kenneth Forbus, Mark Keane, Agnar Aamodt, and Ian Watson. Retrieval, reuse, revision and retention in case-based reasoning. *The Knowledge Engineering Review*, 20(3):215–240, September 2005.

[23] John F. Nash. Non-cooperative games. *Annals of Mathematics*, 2(54):286–295, 1951.

[24] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, New Jersey, 1944.

[25] Maria Nilsson, Joeri van Laere, Tom Ziemke, Peter Berggren, and Birgitta Kylesten. A user study of the Impact matrix, a fusion based decision support for enhanced situation awareness. In *Proceedings of the 11th International Conference on Information Fusion (FUSION 2008)*, Cologne, Germany, June 30–July 3, 2008.

[26] Per-Arne Persson and James M. Nyce. Intuitive tools? Design lessons from the military intelligence community. *American Intelligence Journal*, 25(1):38–50, Summer 2007.

[27] Per-Arne Persson and James M. Nyce. Integrating human effort and technology in the ISTAR model: An ethnographic perspective. Manuscript, National Defence College, Stockholm, Sweden, 2008.

[28] Marco A. Pravia, Ravi K. Prasanth, Pablo O. Arambel, Candy Sidner, and Chee-Yee Chong. Generation of a fundamental data set for hard/soft information fusion. In *Proceedings of the 11th International Conference on Information Fusion (FUSION 2008)*, Cologne, Germany, June 30–July 3, 2008.

[29] Howard Raiffa. *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Addison–Wesley, Reading, Massachusetts, 1968.

[30] Maria Riveiro. Evaluation of uncertainty visualization techniques for information fusion. In *Proceedings of the 10th International Conference on Information Fusion (FUSION 2007)*, Québec, Canada, July 2007.

[31] Thomas C. Schelling. *The Strategy of Conflict*. Harvard University Press, Cambridge, Massachusetts, 1960.

[32] Bruce Schneier. Attack trees. *Dr. Dobb's Journal*, 24(12):21–22, 24, 26, 28–29, December 1999.

[33] Thomas B. Sheridan and William L. Verplank. Human and computer control of undersea teleoperators. Technical report, Man-Machine Systems Laboratory, Department of Mechanical Engineering, Massachusetts Institute of Technology, July 1978.

[34] Hedvig Sidenbladh, Pontus Svenson, and Johan Schubert. Comparing future situation pictures. In *Proceedings of the Eighth International Conference on Information Fusion (FUSION 2005)*, pages 963–968, Philadelphia, Pennsylvania, July 2005.

[35] Mark Steyvers and Tom Griffiths. Probabilistic topic models. In Thomas K. Landauer, Danielle S. McNamara, Simon Dennis, and Walter Kintsch, editors, *Handbook of Latent Semantic Analysis*, University of Colorado Institute of Cognitive Science Series, chapter 21, pages 427–448. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 2007.

[36] Charles Sutton, Clayton Morrison, Paul R. Cohen, Joshua Moody, and Jafar Adibi. A Bayesian blackboard for information fusion. In *Proceedings of the Seventh International Conference on Information Fusion (FUSION 2004)*, volume 2, pages 1111–1116, Stockholm, Sweden, June 28–July 1, 2004.

[37] Robert Suzić and Pontus Svenson. Capabilities-based plan recognition. In *Proceedings of the Ninth International Conference on Information Fusion (FUSION 2006)*, Florence, Italy, July 2006.

[38] Pontus Svenson. Capabilities-based force aggregation using random sets. In *Proceedings of the Eighth International Conference on Information Fusion (FUSION 2005)*, pages 872–878, Philadelphia, Pennsylvania, July 2005.

[39] Pontus Svenson, Tomas Berg, Pontus Hörling, Michael Malm, and Christian Mårtenson. Using the impact matrix for predictive situational awareness. In *Proceedings of the Tenth International Conference on Information Fusion (FUSION 2007)*, Québec, Canada, July 2007.

[40] Swedish Armed Forces. *Grundsyn Underrättelsetjänst* [Intelligence Service Basic View]. M7739-350003, Stockholm, Sweden, 2008.

[41] Gheorghe Tecuci, Mihai Boicu, Cindy Ayers, and David Cammons. Personal cognitive assistants for military intelligence analysis: Mixed-initiative learning, tutoring, and problem solving. In *Proceedings of the First International Conference on Intelligence Analysis*, McLean, Virginia, May 2005.

[42] Victoria Uren, Philipp Cimiano, José Iria, Siegfried Handschuh, Maria Vargas-Vera, Enrico Motta, and Fabio Ciravegna. Semantic annotation for knowledge management: Requirements and a survey of the state of the art. *Web Semantics: Science, Services and Agents on the World Wide Web*, 4(1):14–28, January 2006.

[43] Stanley Wasserman and Katherine Faust. *Social Network Analysis: Methods and Applications.* Structural Analysis in the Social Sciences. Cambridge University Press, New York, 1994.

[44] Franklin E. White, Jr. A model for data fusion. In *Proceedings of the First National Symposium on Sensor Fusion*, volume 2, Orlando, Florida, April 1988.