Capturing Analytic Thought

John D. Lowrance
Ian W. Harrison
Andres C. Rodriguez
Artificial Intelligence Center
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94404
{lowrance, harrison, rodriguez} @ai.sri.com

Abstract
The survival of an enterprise often rests upon its ability to make correct and timely decisions, despite the complexity and uncertainty of the environment. Because of the difficulty of employing and scaling formal methods in this context, decision makers typically resort to informal methods, sacrificing structure and rigor. We are developing a new methodology that retains the ease of use, the familiarity, and (some of) the free-form nature of informal methods, while benefiting from the rigor, structure, and potential for automation characteristic of formal methods. Our approach records analysts' thinking in a corporate knowledge base consisting of *structured arguments.* The foundation of this knowledge base is an ontology of arguments that includes two main types of formal objects: *argument templates* and *arguments.* An argument template records an analytic method as a hierarchically structured set of interrelated questions, and an argument instantiates an argument template by answering the questions posed relative to a specific situation. This methodology emphasizes the use of simple inference structures as the foundation of its argument templates, making it possible for analysts to independently author new templates. When authoring an argument template, the analyst can choose to embed *discovery tools*, which are recommended methods of acquiring information pertaining to the questions posed. An analyst wanting to record an argument selects an appropriate template, uses the discovery tools to retrieve potentially relevant information, selects that information to retain as *evidence* and records its *relevance*, *answers* the questions, and records the *rationale* for the answers. The result is a recorded line of reasoning that breaks down the problem, bottoming out at the documents and other forms of information that were used as evidence to support the answers. The resulting collection of arguments and templates constitutes a corporate memory of analytic thought that can be directly exploited by analysts or automated methods.

K-CAP'01, October 22-23, 2001, Victoria, British Columbia, Canada.

Keywords
Structured arguments, evidential reasoning, analysis, knowledge management.

INTRODUCTION
Understanding the world and facing the different alternatives it presents to us is crucial in any effort. Different studies and formalisms of *argumentation* have come out of different fields such as philosophy [11, 14, 15, 19] decision analysis [17] and artificial intelligence [9, 16, 4]. These formalisms attempt to deal with the uncertainty inherently present in the world. Behind every decision, though, there is an argument supporting it, and arguments range from rhetorical explanations to mathematical proofs. Argumentation theory leverages problem solving under uncertainty by supporting qualitative and quantitative approaches.

Analysis, on the other hand, deals with the examination and separation of a complex situation, its elements, and its relationships. More often than not, the situation is full of unknowns, uncertainties, and deliberate misinformation. The analyst is confronted not only with the facts, but also with his or her knowledge about the facts and assumptions, others' possible knowledge, the hypotheses that can be drawn from those facts, and the evidence supporting and contradicting those hypotheses (Heuer 1999).

Under the sponsorship of the Defense Advanced Research Projects Agency (DARPA) of the U.S. Department of Defense, SRI International is developing SEAS, the Structured Evidential Argumentation System also known as the SRI Early Alert System (Lowrance, Harrison, and Rodriguez 2000). This work builds upon an earlier effort (Stokke et al. 1994) that developed the first SEAS prototype applied to the problem of early warning for project management. In our current work, SEAS is being generalized and applied to the problem of crisis warning for national security. Our goal is to construct a system capable of aiding intelligence analysts in leveraging analytic products and methods developed for past situations or by other analysts addressing the same or similar contemporary problems. These analytic products take the form of arguments: given a framework of assumptions, some conclusions or statements can be reached. While national security analysis is the focus of this work, we believe that the tools and methods being developed have broad application outside of the national security arena. We believe that these tools and methods can be effectively applied to any problem where regular assessments must be made, based upon evidence from multiple sources, within a complex and uncertain environment.

CAPTURING ANALYTIC METHODS

Our approach is based on the concept of a *structured argument.* A structured argument is based on a hierarchically organized set of questions (a tree) that is used to assess whether an opportunity or threat of a given type is imminent. This hierarchy of questions is called the *argument's template* (as opposed to the *argument*, which is an instantiation of the template). This hierarchy of questions supporting questions may go a few levels deep before bottoming out in questions that must be directly assessed and answered. These are multiple-choice questions, with the different answers corresponding to discrete points or subintervals along a continuous scale, with one end of the scale representing strong support for a particular type of opportunity or threat and the other end representing strong refutation. Leaf nodes represent primitive questions, and internal nodes represent derivative questions. The links represent support relationships among the questions. A derivative question is supported by all the derivative and primitive questions below it. Figure 1 illustrates a seventeen-question argument template, with twelve primitive questions and five derivative questions. Note that question 1 is answered based upon the answers to 1.1, 1.2, 1.3, and 1.4, and 1.2 is answered based upon the answers to 1.2.1, 1.2.2, and 1.2.3.

Figure 1: An example argument skeleton

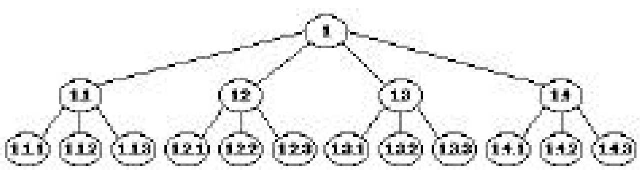

An inference method completes an argument template. It is used to automatically answer some questions based upon the answers to other questions. The analyst answers the primitive questions in the question hierarchy, and the answers to the derivative questions are automatically calculated. In so doing, our approach emphasizes the use of simple and regular inference structures. These structures are captured by *argument skeletons* and associated *inference methods.* The same argument skeleton and inference methods are typically used to support multiple argument templates over widely differing topics. A typical inference method might take the maximum answer as the conclusion when combining several questions assessed along a continuous scale. The idea is that if the argument template author fully understands the structure of the interrelated questions that constitute the argument skeleton and the propagation scheme implemented by the inference method, then the author can write the argument template questions and answers to fit. The simpler the argument skeletons and inference methods, the easier it is for the author.

The use of regular argument skeletons is encouraged – that is, skeletal trees where all branches are identically structured. Regular structures help to encourage that equal time and emphasis are placed on all aspects of an analysis. Likewise, the use of uniform or regular inference methods is encouraged. A uniform inference method, where every derivative question's answer is derived using the same fusion method, makes for the easiest arguments to understand and lines of reasoning to follow. A regular inference method, one that employs the same fusion method across all questions at the same depth in the skeletal tree, is the next easiest to understand and follow.

Our philosophy is directly opposed to that of most uncertain reasoning systems. In most systems, the author begins by determining what questions might be asked and then interrelates them through a complex set of interconnections, typically annotated with conditional probabilities. As a result, the updating scheme is often complex and difficult to follow for those not versed in probability theory. While this "strong model" approach can be very effective when properly applied, we believe that the "weak model" approach emphasized here is easier to understand and use. Its effectiveness is directly related to the author's ability to adapt to these simple and regular inference structures, writing questions and answers that properly function within these constraints. Thus, knowledge is entered via text editing, without the use of probabilities or weights, making knowledge entry easy.

The challenge in authoring an argument template is to break the problem down into a hierarchically structured set of questions that matches the selected argument skeleton and whose interrelationships among the answers follow the inference method. Therefore, it is critical that the author understands the structure of the argument skeleton and the effect of the inference method, before beginning to fashion the questions and answers that will be posed by the argument template. See Figure 2 for an example argument template question hierarchy.

Each derivative question is represented by two text strings: a *topic* and the *question* itself. Primitive questions also include a question *amplification* string and five *multiple-choice answers.* The amplification states the question in more detail, reminding the user of the range of things to consider when answering the question.

To facilitate the rapid comprehension of arguments, we use a traffic light metaphor; relating answers to colored lights along a linear scale, from green to red. The questions in a template are yes/no or true/false; the multiple-choice answers for primitive questions partition this range, associating an answer with each colored light. Typically, a five-light scale is used (green, yellow-green, yellow, orange, red). Here green might correspond to true, red to false, and the other three to varying degrees of

Figure 2: An example argument template question hierarchy

1. **POLITICAL**: Is this country headed for a political crisis?
   - 1.1. **POLITICAL INSTABILITY**: Is political instability increasing?
     - 1.1.1 **INCREASINGLY UNSTABLE/WEAK GOVERNMENT**: Is the government becoming increasingly unstable or weak?
     - 1.1.2 **INCREASING CONFLICT OVER POLICY/ISSUE AREA**: Is increasing conflict over policy/issue areas having a destabilizing effect?
     - 1.1.3 **DECREASING PUBLIC CONFIDENCE**: Is decreasing public confidence in the leadership or government policies having a destabilizing effect?
   - 1.2. **POWER STRUGGLE**: Is there a government power struggle with potentially destabilizing consequences?
     - 1.2.1. **FACTIONALISM**: Is there evidence of growing factionalism within the government, bureaucracy, or legislature that is leading to or exacerbating a power struggle?
     - 1.2.2. **OPPOSITION CHALLENGE**: Is there a significant political opposition challenge to the government that is leading to or exacerbating a power struggle?
     - 1.2.3. **SUBNATIONAL GROUP INFLUENCE**: Are powerful subnational groups contributing to a government power struggle by influencing or backing specific government officials/factions?
   - 1.3. **GOVERNMENT RESPONSE TO SOCIO-POLITICAL DISCORD**: Is the government resorting to increasingly stringent measures in response to socio-political discord with potentially destabilizing consequences?
     - 1.3.1. **REPRESSION OF POLITICAL OPPOSITION**: Is government repression of the political opposition or dissident groups occurring/increasing?
     - 1.3.1 **REPRESSION OF SOCIAL/RELIGIOUS GROUPS**: Is government repression of social/religious groups occurring/increasing?
     - 1.3.2 **NTERNAL SECURITY MEASURES**: Is the government instituting or strengthening internal security measures in response to armed (guerrilla/insurgency/separatist) movements or terrorist/criminal activity?
   - 1.4. **STRUCTURAL/INSTITUTIONAL PROBLEMS**: Are there serious or worsening or institutional problems that could have destabilizing consequences?
     - 1.4.1. **CONSTITUTIONAL CONFLICT/CRISIS**: Is there a constitutional conflict/crisis?
     - 1.4.2. **ERODING LEGAL AUTHORITY/ADMINISTRATIVE FUNCTIONS**: Are legal authorities or administrative functions eroding?
     - 1.4.3. **DECREASING PUBLIC CONFIDENCE**: Is public confidence in public institutions decreasing?

certainty. Ideally, the multiple-choice answers are as concrete as possible and directly and unambiguously observable, making it easier for the user to recognize the answer that fits the situation being analyzed. No multiple-choice answers are associated with derivative questions; within arguments, their answers are strictly summarized by lights indicating their degree of certainty.

There are two distinct ways of approaching the structuring of an argument template: top-down and bottom-up. Using the top-down approach, one starts with the central question and attempts to break it down into a small set of supporting questions, each of approximately the same significance; then one breaks down each of those questions, attempting to break each into the same number of equally significant questions. This procedure continues until questions are produced that can be directly answered or until the number of overall questions has become too numerous to include in a single template. In this latter case, the author might elect to limit the depth of the original template and then capture those elements that fell below that depth limit in their own templates; each of these *cascaded templates* would share its root question with one of the primitive questions in the original template. The relationship of these cascaded templates to the original template can be captured by adding these to the original template as *discovery tools* (more on this below). As such, an analyst who is developing an argument based upon the original template, and is confronted with one of its primitive questions, can either elect to directly answer the stated question or invoke one of these discovery tools to further break down the question. The advantage of this approach is that the analyst determines which of these discovery tools to employ, thus choosing where and where not to spend time.

Using the bottom-up approach, one starts by enumerating the detailed conditions that should lead to warning. Once these are enumerated, one begins to cluster these into coherent collections of roughly equal size and significance. One then clusters the clusters, again striving for clusters of equal size and significance, and continues this process until a single cluster remains. Each cluster should give rise to a question in the resulting template, with the nesting of the clusters captured as supporting questions.

In practice, neither the top-down nor bottom-up approach is employed in its pure form. Instead, both are typically employed at different times, one after the other, until a satisfactory result is achieved. Once the overall skeletal structure has been established, then the author's attention should turn to writing the detailed questions and answers for the template.

In general, *discovery tools* are recommended methods for acquiring information relevant to answering questions in an argument template. These might be links to Web pages, queries to databases or search engines, parameterized launches of other analytic tools, or references to cascaded templates. They capture an important aspect of an analyst's knowledge, namely, where and how to go about seeking information relevant to answering questions. Knowledge of this form is one thing that distinguishes an expert from a novice analyst. Discovery tools are captured on primitive questions within a template by storing their URLs along with short citation strings used to reference them. Again, simple text editing is all that is needed to define these.

Finally, the author should establish a *situation descriptor*, for a new template, that describes the type of situations for which the template is intended to be used. Unlike the other information provided by the user in defining a template, much of the information in a situation descriptor is chosen from a situation ontology rather than being free text. The situation ontology serves much the same purpose as a card catalog in a library; it establishes indices and terms that are useful for retrieving objects based upon the type of situation to which they are applied. For national security problems, these include the part of the world being analyzed (e.g., the continent, region, or country under assessment), the principal actor (e.g., the leadership, the government, or its people), the event (e.g., political, economic, financial, or currency), and the time period. These descriptions, with the exception of time, are selected from hierarchies of terms that are established through traditional knowledge engineering techniques. By indexing objects according to this situation ontology, both exact and semantically close matches can be automatically retrieved based upon a description of the situation of interest expressed in the same terms. These situation descriptors are augmented by free text fields where the specific aspects of the situation can be fully expressed; thus, the ontological terminology need not fully capture every distinction.

In practice, we have found that analysts are capable of authoring templates after minimal training, but that authoring high-quality templates is challenging and requires additional experience. To jump-start this process for problems of national defense, we convened a multidisciplinary team of experts to establish high-level templates for assessing the stability of nation states. The idea was to provide analysts with an example that they could then improve upon or adapt to specific situations, because it is easier to modify an example than to generate a template anew. The results have been well received. We imagine that variants of this high-level template will eventually be supported by cascaded templates that are more pointed. While the high-level template is useful in reminding analysts of the full range of indicators that need to be assessed and for generally organizing the analysis, their abstract nature prevents them from delivering much in the way of expert guidance. Given that templates cascaded under this high-level template will address more specific and limited analytic tasks, we anticipate that they will capture expert knowledge suitable for guiding analysts in doing analytic tasks that fall outside of their areas of expertise. Thus, these templates capture and deliver best practice.

CAPTURING ANALYTIC PRODUCTS

Arguments are formed by answering the questions posed by a template and attaching the evidence that was used in arriving at the selected answers. In essence, an argument organizes the indications and warning signs for the given type of opportunity or threat.

Answers are chosen from the multiple choices given by the associated template. If the available information does not allow the analyst to reduce the possible answers to a single choice, multiple ones can be selected bounding the answers that remain possible, given the available information. The *rationale* for answering in that way is recorded as a text string with attribution given to the answering analyst and the time that that answer was given.

Upon answering each question, the template's inference method is applied, deriving the answers to derivative questions. Using the traffic light metaphor, arguments can be displayed as a tree of colored nodes. Nodes represent questions, and colors represent answers. Figure 3 shows one such tree. The line of reasoning can be easily comprehended and the analyst is able to quickly determine which answers are driving the conclusion. By examining the high-value answers, the rationale behind the line of reasoning can be understood.

Figure 3: Argument hierarchy showing answers

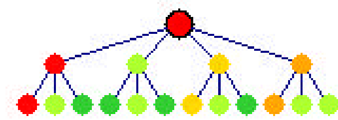

Information used as evidence to support the answers given in an argument is recorded as part of the argument. When information that is potentially relevant to answering a question posed is first found, it is entered as an *exhibit*. An exhibit assigns a unique identifier to the information, and records the URL for accessing it and a *citation* string for referring to it (typically consisting of some combination of title, author, and date). When the *relevance* of the information to the question at hand is determined, the exhibit is promoted to *evidence*. The relevance is recorded in two ways: as a text string explaining the significance and as the answer(s) to the question that would be chosen if the answer were to be based solely upon this evidence. The analyst making this assessment and the time of the assessment are recorded as well. When evidence is present, the rationale typically explains how the collective evidence supports the answer(s) chosen, explaining away that evidence that contradicts the answer and weaving together the supporting evidence to arrive at the stated conclusion.

When discovery tools are present, they can be used to aid in the collection of evidence. When these tools are based upon cascaded templates, cascaded arguments result from their use. In this way, the analyst can choose where they want to do a more thorough analysis, delving more deeply in a targeted way. A cascaded argument's conclusion can be automatically used as its relevance in support of the higher-level argument.

The analyst also chooses a *fusion method* for combining all of the evidence gathered supporting a single template question. The fusion method can be manual (i.e., the analyst answers the question based on his or her understanding of the evidence and its relevance) or automated (i.e., the answer is automatically reached by applying a combination method to the relevance of the supporting evidence). When an automated method is in use, changes in supporting arguments can ripple up through the arguments that they support, changing their conclusions.

As seen in Figure 4, complex lines of reasoning can be captured using this methodology. Here a multidimensional argument (i.e., a coordinated set of unidimensional arguments like those discussed) is graphically depicted at the top; it represents a coordinated assessment along multiple perspectives. It is supported by structured arguments as well as documents and analytic products produced by other tools. This structure allows analysts to quickly come to understand the reasoning of others and compare and contrast it with their own.

Figure 4: Cascaded structured arguments

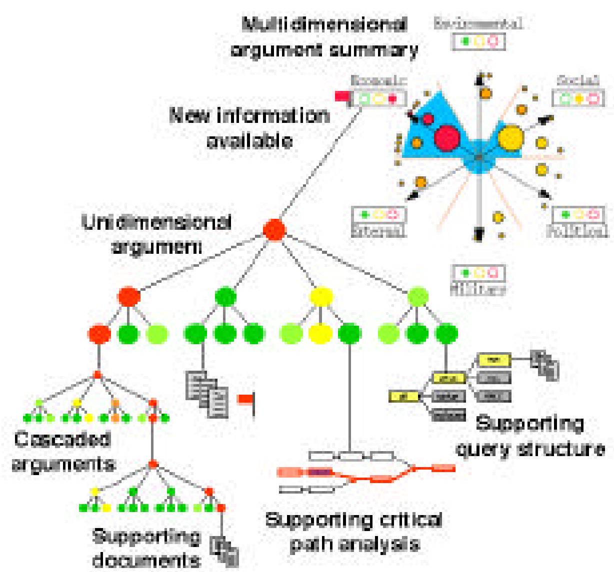

Like argument templates, arguments too have associated situation descriptors. An argument's situation descriptor is like a template's situation descriptor except that it captures information pertaining to the prevailing situation for which the argument was developed. Like the situation descriptors associated with templates, they are used to find arguments that address related situations.

A CORPORATE MEMORY OF ANALYTIC KNOWLEDGE

To support the application of the structured argumentation methodology, SRI is developing SEAS, the Structured Evidential Argumentation System. SEAS has been developed as a Web server that communicates with remote browser-based clients. Through HTML and JavaScript, SEAS supports analysts in locating, understanding, and developing templates and arguments. This analytic knowledge is maintained within a knowledge-base management system, with ephemeral views served up upon demand. Figure 5 shows one such view of a primitive question within an argument.

Figure 5: SEAS argument in browser client

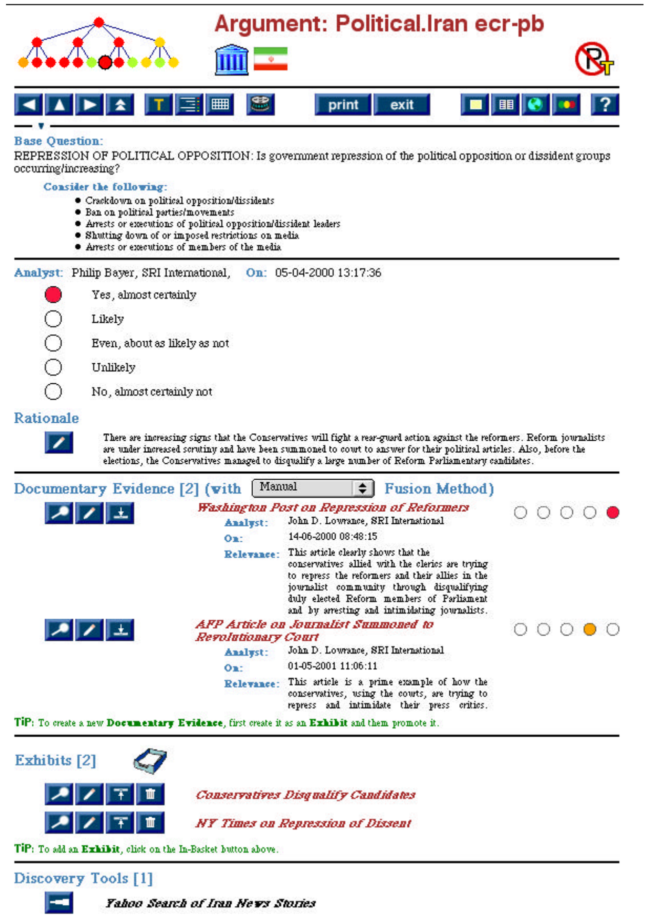

If we are to recognize future opportunities and threats, then we must relate the present to the opportunities and threats of the past. We must understand how the current situation is like or unlike previous situations; how the indications and warning signs are similar or dissimilar; how previous opportunities or threats were recognized or missed; how previous opportunities or threats evolved and thereby how the present situation might evolve; and how previous situations were leveraged, mitigated, exacerbated, or missed. In short, we need a corporate memory that is more than a historical data repository; we need a corporate memory of analytic products and methods on which to base future analysis.

By recording and retaining analytic thinking in a common knowledge repository, analysts can leverage the thinking from the past and present when addressing new tasks. Based upon the indexing provided by the situation descriptors, potentially relevant templates and arguments can be found.

Beyond the analytic methods (i.e., argument templates), analytic products (i.e., arguments), and their associated situations (i.e., situation descriptors), we have found that analysts need additional means for associating meta-knowledge with these objects. To address this need, SEAS supports *memos.*

Memos are structured annotations that are attached to other objects within the SEAS knowledge base. Each memo includes text strings for its *subject* and *body* and a *type* selected from a pre-established set including critique, to do, summary, instruction, and assumption. Like arguments and templates, they have a designated *audience* that restricts their access by others; only those that are members of the audience will know of their existence. As such, memos provide a means for private, semiprivate, or public communication among analysts. Critiques are a way for contemporary analysts to contribute to each other's work. Assumptions might be added so that analysts in the future will better be able to interpret a historical analysis. Within SEAS, memos can be selectively filtered based upon their type, with graphical depictions indicating to the user where they can be found. This provides a ready means for analysts to find and interpret this form of meta-knowledge.

While analytic knowledge that is developed in SEAS is retained in its corporate memory, as are references to external analytic products used as evidence, there are times when one would like to import arguments produced using other technologies, so that they can be extended or otherwise modified. Our objective is to provide a means for the exchange of information among tools that can be said to produce arguments. If tools can be said to be argumentation tools, then they should be able to exchange arguments. Although argumentation tools share common concepts, they invariably have some unshared concepts, necessarily making importation imperfect.

Toward this objective, we are defining the Argument Markup Language (AML), an XML representation of arguments, and modifying SEAS to support the importation and exportation of these objects. The initial set of argumentation tools that we aim to support comprises those based upon Bayesian nets, particularly drawing from the Bayesian Net Interchange Format (Microsoft 2001), CIM (Veridian 2001), a structured argumentation tool developed at the same time as SEAS but with an emphasis on arguments about processes, and SEAS. While this is the initial set, we are aiming for a general design that will support a far greater number of tools, including those based upon both numeric and symbolic representations of certainty. We began by looking for common semantic concepts within these tools and using terminology from the Law to capture them. Legal terminology was selected since the Law already includes a rich notion of argumentation from evidence and provides a technology-neutral vocabulary, many of whose terms are in common use. An initial version of AML has been defined, and CIM and SEAS are being modified to support it.

RELATIONSHIP TO OTHER WORK

The structured argumentation methodology and SEAS were developed to aid those performing analytic tasks. In particular, we were not looking to automate the analytical reasoning that they perform, but to facilitate it. This methodology

- Encourages careful analysis, by reminding the analyst of the full spectrum of indicators to be considered
- Eases argument comprehension and communication by allowing multiple visualizations of the data at different levels of abstraction, while still allowing the analyst or decision maker to "drill down" along the component lines of reasoning to discover the detailed basis and rationale of others' arguments
- Invites and facilitates argument comparison by framing arguments within common structures

In addition, SEAS provides synchronous and asynchronous access to a corporate memory of analytic methods and results, which allows analysts to work together on common arguments as well as leveraging historical results. Collaboration, then, is recognized as an important part of the process and leads to arguments that are richer than would have been otherwise the case. The Web is an ideal medium for collaboration, driven by the near ubiquity of browser software and the information explosion on the Web.

The goals of structured argumentation differ from those of other knowledge capturing tasks. In most knowledge engineering efforts the objective is to elicit and represent the knowledge of humans in machines so that the machine can later use this knowledge to approximate the reasoning of humans. This largely requires that the knowledge be captured in not natural language but in ontological structures that can be more readily manipulated by machines. Examples include Cyc [1], DARPA High Performance Knowledge Bases [2], DARPA Rapid Knowledge Formation [3], GKB-Editor [13], EcoCyc [8], and Ontolingua [5].

Today, intelligence analysts usually capture their knowledge in text documents. Typically, these documents have minimal structure, limited to section titles that break up the document. These intelligence reports are intended for human consumption. However, because of their limited structure they are time consuming to read and understand. To compare one report with another requires that both reports be read, and it is up to the reader to find common and uncommon aspects of the underlying reasoning. It is also up to the reader to extract the analytic method if it is to be employed in doing related analyses. Searching a collection of such reports to find ones that might be related to the current problem of interest is also time consuming. Of course, word processing and search engines can help to speed this process, but the level of aid is fundamentally limited.

Structured argumentation fits between these two approaches. It introduces more structure into the analytic environment than is in use today but not as much as typical knowledge engineering efforts. The analytic method is separated from the analytic products, resulting from its application. The analytic method is broken down into a set of smaller analytic tasks, with their interrelationships captured. Methods for acquiring information in support of these analytic tasks are also broken out. In structure, analytic results parallel the analytic methods on which they are based, with links to the information that supports the conclusions retained, and to the interpretations of that information relative to each analytic task. The type of situation for which a method was designed and for which a result was produced is also captured. However, much of the knowledge captured remains in natural language. In fact, when one compares an analytic product produced using SEAS with a contemporary analytic product expressed in a text document, one finds that most of the text found in the document is found within the structured argument. The structure has not replaced the words as much as it has augmented them, making it possible for the machine to aid the analysts in new ways.

This approach bears a resemblance to some recent work in support of knowledge mobility [6]. In this work, the rationale and sources of knowledge, drawn upon in engineering knowledge, are retained within the resulting formal structures. So doing allows others to more readily understand why knowledge was captured as it was, making it easier to reapply, extend, or modify. These resilient hyper knowledge bases use a layered architecture to capture knowledge from the most to the least formal. This work can be viewed as building up from formal knowledge to less formal knowledge, while our work can be viewed as building down from informal knowledge to more formal knowledge.

The structure introduced into the analytic process by the structured argumentation methodology, although motivated by the desire to help humans, also represents an opportunity for greater accessibility to automated methods. These methods might attempt to provide critical feedback to the analyst or automatically make corrections. Such feedback can be readily communicated using the SEAS memo facility; thus, automated collaborators would interact in the same way as human collaborators. Some of these capabilities could be introduced without the need to perform any natural language understanding; other capabilities might require some limited understanding; still others would benefit from more comprehensive natural language understanding.

Without the introduction of natural language understanding, we intend to develop an automated argument critic that provides several kinds of feedback. For example, such a critic could examine the answer to every question in an argument, to determine if the answer is supported by evidence, if each piece of evidence includes a statement of its relevance, and if the rationale for the overall answer is given. It could also check for overreliance on any single document supporting the answer to multiple questions, since overuse of any source of information leaves one vulnerable to its accuracy and truthfulness.

The corporate memory of arguments presents other opportunities. By comparing the focal argument to successful arguments from the past, other useful sources could be identified that have not been used in the focal argument. Likewise, sources that had previously led to poor results could be flagged. Similarly, more complex patterns of previous use could be exploited.

With the aid of some natural language understanding technology, coupled with inference capabilities based upon formal knowledge representations, we might develop more sophisticated aids. These might look to find logical contradictions in the way that evidence was interpreted or in the rationale accompanying answers given. They might also look to suggest alternative interpretations.

CONCLUSIONS

We believe that our structured argumentation methodology, as implemented in SEAS, has shown that the addition of even minimal structure into the analytic process can aid analysts in developing, communicating, explaining, and comparing analytic results. An important aspect of this methodology is the retention of direct links to the source material and its interpretation relative to the conclusions drawn, allowing analysts to readily comprehend the thinking of others. This, coupled with a collaborative environment and a corporate memory of analytic thought, retaining the analytic methods and products of an enterprise, allows analysts to leverage the thinking of others both past and present. Finally, even though our methodology was motivated by the desire to help human analysts, it lays the groundwork for the introduction of automated methods to substantially aid or partially supplant human analytic reasoning. We contend that this methodology complements those knowledge capturing methodologies that strive to formally represent human knowledge in rich ontological structures.

ACKNOWLEDGEMENTS

We would like to recognize the contributions made to this work by our colleagues supporting DARPA's Genoa program. In particular, we would like to thank the present and past DARPA Genoa program (Doug Dyer, Tom Armour, and Brian Sharkey), the other Genoa contractors, particularly Syntek (John Poindexter and Greg Mack), Veridian Systems Division (Mark Lazaroff, Scott Fisher, and Jill Jermano), and ISX (Mark Hoffman), and our colleagues at SRI (Tom Boyce, Eric Rickard, Phil Bayer, Vinay Chaudhri, Jerome Thomere, and Karen Myers).

REFERENCES

[1] Cycorp (2001); The Cyc Knowledge Server; http://www.cyc.com/

[2] DARPA (2000); High Performance Knowledge Bases Project, http://reliant.teknowledge.com/HPKB/

[3] DARPA (2001); The Rapid Knowledge Formation Project; http://reliant.teknowledge.com/RKF/

[4] Dung, P. (1995); On the Acceptability of Arguments and Its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games; *Artificial Intelligence* 77 (pp. 321-58)

[5] Fikes, R., Farquhar, A., and Rice, J. (1997); Tools for Assembling Modular Ontologies in Ontolingua; Knowledge Systems Larobatory, Stanford University

[6] Gil, Y. (2001); Knowledge Mobility: Semantics for the Web as a White Knight for Knowledge-Based Systems, forthcoming

[7] Heuer, R. (1999); *Psychology of Intelligence Analysis*; Center for the Study of Intelligence, Central Intelligence Agency

[8] Karp, P., Chaudhri. V., and Paley, S. (1999); A Collaborative Environment for Authoring Large Knowledge Bases; in *Journal of Intelligent Information Systems*, Vol. 13 (pp. 155-194)

[9] Loui, R. (1987); Defeat Among Arguments: A System of Defeasible Inference; in *Computational Intelligence*; Vol 3 (pp. 100-106).

[10] Lowrance, J., Harrison, I., and Rodriguez, A. (2000); Structured Argumentation for Analysis; in *Proc. 12th Int. Conf. on Systems Research, Informatics, and Cybernetics: Focus Symposium on Advances in Computer-Based and Web-Based Collaborative Systems* (pp. 47-57)

[11] Lorenzen, P. and Lorenz, K. (1977); *Dialogische Logik*; Wissenschaftliche Buchgesellschaft Darmstadt

[12] Microsoft (2001); XML Belief Net File Format; http://www.research.microsoft.com/dtas/bnformat/

[13] Paley, S., Lowrance, J., and Karp, P. (1997); A Generic Knowledge Base Browser and Editor; in *Proc Ninth Conf. on Innovative Applications of Artificial Intelligence*

[14] Perelman, C. (1970); *Le Champ de l'argumentation*; Bruxelles: Éditions de l'Université

[15] Perelman, C. and L. Olbrechts-Tyteca. (1958); *Traité de l'argumentation - la nouvelle rhétorique*; Bruxelles: Éditions de l'Université

[16] Pollock, J (1987); Defeasible Reasoning; in *Cognitive Science*, Vol. 11 (pp. 481-518).

[17] Sycara, K. (1990); Persuasive Argumentation in Negotiation, *Theory and Decision* Vol. 28, No. 3 (pp. 203-42)

[18] Stokke, R., Boyce, T., Lowrance, J., and Ralston, W. (1994); Evidential Reasoning and Project Early Warning Systems; *Journal of Research and Technology Management*

[19] Toulmin, S. (1958); *The Uses of Arguments*; Cambridge University Press

[20] Veridian Systems Division (2001); Critical Intent Modeling; unpublished