StructuredArgumentationforAnalysis

J. D. Lowrance, I. W. Harrison, A. C. Rodriguez
Artificial Intelligence Center
SRI International
333 Ravenswood Avenue, Menlo Park, CA 94025, USA
{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 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, 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 aims to foster thoughtful and timely analysis through the introduction of structure, and collaboration through access to current and past analytic results. By recording analysts' thinking in structured arguments, the results are more comprehensive, more readily understood, and more easily compared. By providing access to a corporate memory of analytic methods and results, analysts can work together on common arguments and leverage historical results. The structured argumentation methodology encourages a careful analysis, by reminding the analyst of the full spectrum of indicators to be considered. It also eases argument comprehension by allowing the analyst or decision maker to "drill down" along the component lines of reasoning to discover the basis and rationale of others' arguments, and it invites and facilitates argument comparison by framing arguments within common structures.

Introduction

Understanding the world and facing the different alternatives it presents us is crucial in any effort. Different studies and formalisms of *argumentation* have come out of different fields such as philosophy (Lorenzen and Lorenz 1977; Perelman (1970); Perelman and Olbrechts-Tyteca 1958; Toulmin 1958), decision analysis (Sycara 1990), and artificial intelligence (Dung 1995). 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 Project Agency (DARPA) of the U.S. Department of Defense, SRI International is developing SEAS, the Structured Evidential Argumentation System. 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.

Near the inception of this effort, we met with representatives of the intelligence community to discuss the methods currently employed in the production and dissemination of intelligence products and the problems previously encountered when attempts were made to introduce automated tools into this community. Today, analytic products are recorded as text, be they messages, Web pages, or documents. The drawback with this form of recorded arguments is that they are time consuming to read and comprehend, slowing comparisons and explanations. Although text-processing tools help to streamline production of these products, they provide little or no support in guaranteeing that the arguments are comprehensive and easily understood. Many of the attempts made to introduce automation into the intelligence community, were based on decision theoretic constructs and tended to suffer from the same three drawbacks:

- These models reduced the analyst's role to that of data entry, making them unpopular with analysts
- The explanations given for the conclusions were often incomprehensible to analysts, since the lines of reasoning were explained in terms of conditional probabilities, making the conclusions hard to accept
- Even when an analyst came to accept a role limited to data entry and came to believe the conclusions based on a history of success, the world would change, requiring that the analytic model be updated, which the analysts were unable to do

Thus, the tools typically fell into disuse even when they had been initially successful. One of our foci was to develop a user-centric tool. We want users of different levels of expertise to be able to engage in the process of argument creation and analysis. Argument templates will help bridge the gap between an expert and a novice user. A situation ontology will set the range of phenomena an organization wants to cover.

Basic Framework

SEAS is based on the concept of a *structured argument.* A structured argument is based on a hierarchically organized set of questions (a tree) that are 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. Derivative questions are supported by all the derivative and primitive questions linked below it. Figure 1 illustrates a six-question argument template, with four primitive questions and two derivative questions. Note that question `Q1.2` is answered by using the answers to `Q1.2.1` and `Q1.2.2`.

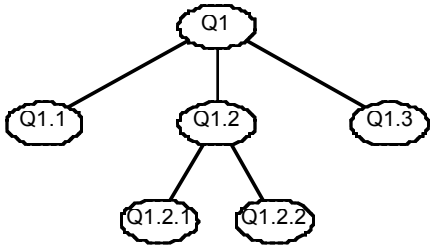

Figure 1: An example argument template

An inference method completes an argument template. SEAS automatically answers some questions based upon the answers to other questions. The analyst answers the leaf questions in the question hierarchy, and SEAS automatically answers the other questions. In so doing, SEAS 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 to understand.

The challenge in authoring a SEAS 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.

Arguments are formed by answering the questions posed by a template and attaching the evidence that was used in arriving at the selected answers. While the argument is under construction, SEAS readily highlights the information that needs or remains to be collected to confirm or deny this hypothesis. In essence, such an argument organizes the indications and warning signs for the given type of opportunity or threat.

An argument creator chooses a *fusion* method for all the evidence gathered. The fusion method can be manual (the author will answer the question based on his or her understanding of the evidence and its particular relevance) or automated (the answer is reached by some combination of the relevance of the evidences using the fusion method associated with the question).

The primary function of SEAS is to remind the analyst of the key indications and warning signs and to record the analyst's reasoning in applying an argument template to the target situation, all to determine the applicability of the given type of opportunity or threat to this situation. The resulting instantiated argument is then ready for other analysts to browse and critique.

The SEAS 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 by SEAS 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.

SEAS also differs from other uncertain reasoning systems in emphasizing the use of small and shallow question hierarchies. This is partially motivated by the desire to make the arguments easily understood in a short amount of time. The more important motivation is the desire to allow the analyst to determine how much time should be spent on different aspects of the problem. That is, rather than basing an argument on a large, comprehensive argument template that covers all aspects of the problem, SEAS emphasizes the use of a collection of cascaded templates, where the leaf question in one argument is answered by the root question in another. In so doing, an analyst can choose to directly answer a leaf question, thus limiting the amount of time spent on that aspect of the problem, or create an argument based upon a cascaded template which will require more time but will provide a more comprehensive analysis of this aspect of the problem. At one extreme an analyst might choose to create only a single high-level argument, using it to record thinking in a structured form; at the other extreme, an analyst might choose to develop a high-level argument through the creation of multiple supporting arguments, helping to guarantee that he or she has done a thorough

analysis of the situation, by taking advantage of detailed analytic methods developed by others more expert in these areas.

To facilitate the rapid comprehension of arguments, SEAS uses a traffic light metaphor; displaying answers to questions in the form of colored lights along a linear scale. By using this metaphor hierarchical arguments can be displayed as a tree of colored nodes. Nodes represent questions, and colors represent answers (red indicates warning, green indicates everything is ok). Figure 2 shows one such tree. The line of reasoning can be quickly comprehended and the analyst is able to look over the entire argument very quickly.

Figure 2: Argument hierarchy showing answers

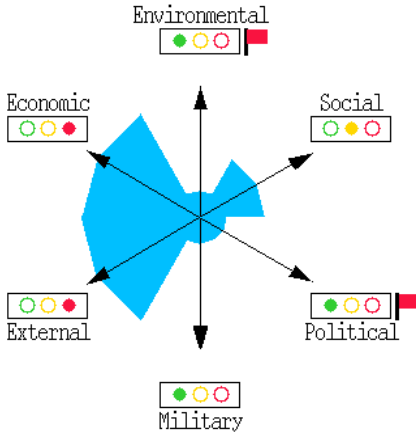

Figure 3: Multidimensional argument view

Effective summaries of full arguments can be printed or displayed on a couple of pages/screens. One technique simply uses indentation of questions to reflect their hierarchical relationship (i.e., outline style) with the colored lights representing the answers on the right. This allows the analyst to "drill down" along an argument's lines of reasoning without leaving the page/screen. Another technique organizes the answers to the uppermost questions in a pattern that resembles spokes on a wheel, as in Figure 3. Each "spoke" corresponds to one answer. Answers are plotted as points along the spokes with the "hub" of the wheel corresponding to the favorable end of the linear scale and the "rim" to the unfavorable. The points plotted on neighboring spokes are connected by lines, and the resulting polygon is filled. The spokes are arranged so that questions falling in common categories are near each other. The result is a plot that visually conveys the argument. This summary is useful for multiperspective arguments (termed m*ultidimensional arguments)*, where the goal is not to produce a single answer, but to provide multiperspective insight on a subject. A multidimensional argument is made up of individual arguments. Analysts can quickly

associate different visual shapes with different argument characteristics. This technique invites comparison when multiple arguments are displayed on a single page.

Using this technique, a single analyst might instantiate several different opportunity or threat models to the same set of data, resulting in multiple arguments that can be compared and contrasted. Multiple analysts might each instantiate the same model and compare their results. The same model might be reapplied to an evolving situation over time to help characterize the significance of changes in the situation. An analyst applying a model is reminded about what missing information is important to an understanding of the situation and can then take steps to acquire it.

Collaboration and Corporate Memory

SEAS seeks to foster collaboration among analysts. In reviewing why intelligence analysts might seek out other analysts, we identified six reasons:

- To learn from others by reviewing their analytic methods and products
- To stimulate creative thinking by rapidly exchanging and generating ideas (i.e., brainstorming)
- To gain insights by having others critique their work
- To share the workload, and thus to get results quicker and to get superior results by having different people do what they do best
- To improve their understanding by comparing and contrasting their results with the results of others
- To improve the quality of their results by combining them with the results of others

Note that most of these activities stress the need for asynchronous collaboration aids. The most important capabilities for supporting collaboration in SEAS are through tools that aid argument or template understanding, argument or template comparison, argument or template merging, and argument or template critiquing, and that support division of labor regarding the creation and editing of arguments or templates.

Collaboration with historical results is also important. 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.

Argument and template understanding and comparison are aided by supporting queries against the corporate memory of arguments and templates to find other analytic products and methods that pertain to similar situations. These queries are based upon an ontology of

situations that serves as the means of cataloging and retrieving SEAS objects according to the time, place, event, and actors involved in situations. Understanding and comparison are also aided by the graphical views and summaries that SEAS produces, permitting retrieved arguments and templates to be rapidly reviewed and to have their differences and similarities understood.

Since SEAS is meant to support a community of analysts, it must address issues of privacy and access. An analyst in the early stages of argument development might not want his or her work to be accessible by others. During development, an analyst might want certain individuals or groups to aid the process by reviewing or contributing to it. Even when an argument is complete, the analyst will want to control who will be allowed to see the results. Further, when an argument is used as documentary evidence in support of another argument, then that argument serving as evidence must be guaranteed to persist in its current state to guarantee the integrity of the argument it supports. To address these issues of access control and referencing in corporate memory, SEAS incorporates the concept of publishing. Three key attributes are related to the two states of publishing: unpublished and published. Published arguments and templates are guaranteed to persist, that is, they will continue to exist; no such guarantee is made for unpublished arguments or templates. As a consequence, only published arguments and templates can be reliably cited, much as only published works are (typically) included in bibliographies so that the reader has a real opportunity to obtain and read them. Unpublished arguments and templates are distinguished from published ones in that they are unstable, that is likely to change in content. Published arguments and templates will not change. Finally, unpublished arguments and templates are distinguished from published ones in that their authors are given write access, while published ones restrict both their authors and audiences to read access.

All arguments and templates originate as unpublished works with a single author. While they remain unpublished, the author can add additional authors. Only the authors have access, and they are free to make modifications as they see fit. It is through this means that an analyst can enlist the help of other analysts in directly contributing to the development of an argument or template. An analyst can indirectly enlist the help of other analysts by citing arguments produced by them as documentary evidence to support their argument, or by making use of templates developed by others as the basis for their arguments. Once a draft argument or template is ready for limited external review, the authors might add people or organizations to the audience. It is risky for this audience to cite this unpublished work since it might go away or be substantially changed in the future. When the authors decide that the argument is ready for external release, they publish it, giving read access to a specified audience in addition to themselves. These published arguments and templates can be reliably cited and referenced in other arguments since they are guaranteed to persist in an unchanging state and constitute the stable basis for corporate memory.

SEAS Use Scenario

The intended use scenario of SEAS is as part of an iterative process of defining and refining an enterprise's view on a particular subject. Our current focus is on the national security arena where an analyst's role might be to develop and communicate an understanding of a country,

region, organization, or technology. In a business setting one could imagine an analyst's role being to develop an understanding of the competitive landscape or the market for a particular product or service.

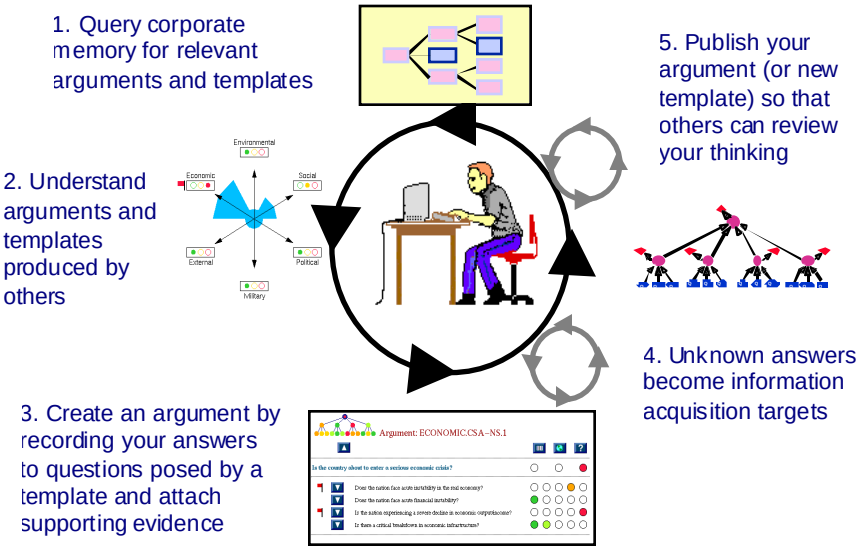

Figure 4: SEAS process cycle

The iterative process supported by SEAS, that an analyst follows when developing an argument, is as follows (see Figure 4):

1. The analyst queries the corporate memory of arguments and templates by describing the current situation, in terms of who, what, where, and when: terms can be chosen from a hierarchical catalog of predefined types. Previously developed arguments and templates are indexed using such meta-data (which we term a *situation descriptor*).
2. Using these criteria the analyst retrieves arguments and templates, in the corporate memory, that have been previously produced by others. These retrieved objects help the analyst understand others thinking regarding this type of situation. At this point the analyst can choose to either create a new argument from scratch using an existing argument template; copy and amend an existing argument; or create or modify a template and use this to create a new argument.
3. The role of an analyst who is developing an argument, is to answer as many questions as possible, and to attach supporting evidence, together with a textual rationale for the way the questions were answered. Unlike some other decision tools, SEAS does not require all answers to be given before it responds. Rather, it continuously represents the current belief in an argument, given the currently supplied information. This belief is updated as questions are answered. The overall belief in the argument, as well as the belief in all the questions in the argument hierarchy, is constantly viewable to the analyst, thus giving

instant feedback. To support answers, an analyst attaches documentary evidence. This can be in the form of a citation to a paper resource, a URL to a resource, or another SEAS argument. While developing the argument an analyst can attach documents that he or she feels might have some relevance to the argument, but which have not yet been analyzed properly. Such evidence, called *exhibits,* can be promoted to documentary evidence at a later date when the analyst has had time to analyze them and determine relevance to the question's answer.

4. Questions that the analyst is unsure how to answer become information acquisition targets. Questions without answers are readily identifiable in the interface, allowing analysts to quickly hone in on parts of an argument that require work. Argument templates, in addition to containing the questions, contain discovery tools. Discovery tools are parameterized queries or tools that the template author believed were relevant methods for acquiring information in answering a particular question. The analyst can fire off these discovery tools, and if any relevant information is found can attach this result as an exhibit to the question. Newly created exhibits are signaled using a mail flag metaphor in the interface.
5. An analyst wanting to share results, or to collaborate with others can use SEAS access control mechanisms to add other authors to the argument (allowing these people to alter the argument), or add others to the audience for review purposes (read-only). When satisfied with an argument, an analyst can publish an argument to the corporate memory. As explained, publishing locks down an argument, guaranteeing that it is both stable and persistent. When publishing, the author is required to describe the situation meta-data (that is to provide a situation descriptor). These descriptors then become the basis for queries against the corporate memory.

Web-based Implementation

SEAS is a web-based client-server application. The client is a traditional Web browser, with the dynamic content being in HTML. The SEAS server, which includes its own HTTP server, generates the dynamic client interface, and processes the client requests. Figure 5 shows a screen shot of the SEAS client for a primitive question.

Our web-based client-server design was motivated by several factors:

- Using a Web browser as the client software almost guarantees the ability for SEAS to be run on any machine. Considering the wide applicability of the SEAS software, the ubiquity of Web browsers, and the heterogeneous computing environments in which SEAS is likely to be employed, it was essential that client software setup be minimized to as close to zero as possible.
- Collaboration is essential to the process of using SEAS. Although much of the need is for asynchronous collaboration, the ability to collaborate synchronously on SEAS arguments and templates was paramount. A web-based client-server architecture, using a browser as the client, gives you such support for free.

- Much of the process in developing SEAS arguments is in discovering and analyzing documentary evidence to answer questions. Increasingly, this evidence comes from websites, and in some cases exists only on the Web, whether on the public Internet itself or on private intranets or extranets (we provide an archiving mechanism to preserve web pages used as evidence). By being a web-based application, SEAS can seamlessly hook into these resources.

The SEAS server also contains an abstract layer for knowledge management, allowing the SEAS server to seamlessly work with different data management systems, where the corporate memory containing the arguments and templates actually resides. This design facilitates the integration of SEAS into enterprise legacy software systems.

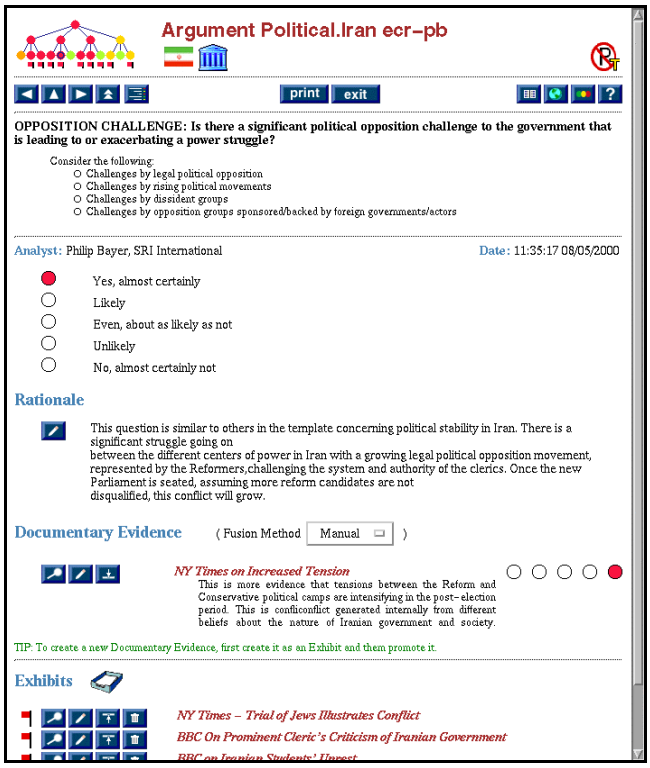

Figure 5: Screen shot of a primitive question in SEAS

Concluding Remarks

In the practical adoption of formal methods of handling uncertainty in decision making, a key barrier to adoption has been the difficulty the end user community has found in developing the required models. With SEAS we have striven to minimize this issue by developing a new methodology, which combines characteristics of both informal and formal methods for handling uncertainty. The structured argumentation methodology provides a framework that has several key benefits:

- 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 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.

From our previous success with the methodology in the domain of early warning for project management, 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. SEAS will soon be taken into an operational setting within parts of the intelligence community, which will allow us to conduct experiments and validate how effective the structured argumentation methodology is in this intelligence domain. Success here would, we believe, transfer across to other areas of business intelligence, such as competitor, market, and technology intelligence. SEAS can provide a structured way for companies to develop and disseminate this kind of knowledge. This will facilitate understanding within an enterprise, and will provide early warning indicators for executives, allowing them to make correct and timely decisions despite the complexity and uncertainty of the environment.

References

Dung, P. M. (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)

Heuer, R. J., Jr. (1999); *Psychology of Intelligence Analysis*; Center for the Study of Intelligence, Central Intelligence Agency

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

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

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

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

Stokke R., Boyce T.A., Lowrance, J.D. and W.K. Ralston, Jr. (1994); Evidential Reasoning and Project Early Warning Systems; *Journal of Research and Technology Management*

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