

Designing a System for Structured Assessment of Compliance Risk

John D. Lowrance

Artificial Intelligence Center
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94404
lowrance@ai.sri.com

Rajesh Ragoobeer

Compliance Lab
IRS Large and Mid-Sized Business Division
1301 Clay Street
Oakland, CA 94612
Rajesh.Ragoobeer@irs.gov

Abstract

The ability to detect, classify and quantify high-risk compliance patterns is crucial in improving compliance risk assessment and formulating effective enforcement strategies within the IRS.

A formal reasoning technique known as structured argumentation has been explored by the IRS LMSB Research to improve compliance risk assessment. The IRS together with SRI International evaluated the use of structured argumentation using SEAS technology, towards enabling a more systematic and innovative approach towards assessing complex compliance issues. Technical experts, auditors, attorneys, researchers and managers can work collaboratively to unravel and respond to high-risk compliance patterns more rapidly and preserve this knowledge to be shared and referenced later.

The Need for a New Perspective

Tax avoidance by larger businesses, wealthy individuals and sophisticated tax practitioners exploit complexity to avoid detection and confuse IRS auditors.

1. The use of multiple filing entities and filing periods to achieve a tax benefit for a specific taxpayer. IRS risk analysis methods that look only at items on a single return are unable to detect complex tax avoidance strategies. A multi-entity, multi-year analysis is required to unravel complex tax avoidance transactions.
2. No data footprint on a tax return to quantify the real compliance risk. Promoters of tax shelters use elaborate strategies that make it hard to link high-risk compliance issues to actual tax return data e.g. the rise in the use of flow-through entities, special purpose entities, complex transactions relationships etc. all help to conceal the tax impact of a tax sheltering device. In testimony before the US Senate Finance Committee testimony (Oct 21, 2003) stated: *Tax shelters having a design or implementation intended to conceal unfavorable facts were often referred to and promoted by certain tax shelter partners as having “good optics.”* IRS needs to better understand the nature and scale of these compliance risk patterns.

3. Previous attempts at improving risk detection and assessment processes focused on automated issue scoring and classification systems. These approaches are not effective at addressing recent tax shelters encountered since current expert knowledge of compliance risk patterns is not considered.
4. The ability to make better use of the specialized technical subject matter experts, the contextual knowledge of revenue agents and data analysis expertise of researchers is of paramount importance in reaching faster and better shared understanding of complex issues.
5. Organizational boundaries can hinder an effective, coordinated risk assessment of complex related return structures and transactions that span multiple IRS divisions (LMSB, SBSE, W&I, TEGE).

A: IRS Test of SEAS¹ Technology:

The Compliance Lab Approach

The Compliance Lab is a research team within the LMSB division of the with the mission of finding newer and more effective methods of analyzing compliance risk to improve workload selection models for flow-through entities.

To test the SEAS technology, the subject of this paper, we worked with revenue agents and technical advisors on an analysis of the “SC²” shelter. SC² was devised and promoted by KPMG and has characteristics typical of other recent schemes promoted to businesses and wealthy individuals²:

1. It involved the use of multiple tax entities over multiple years.

¹ **Note:** SEAS – (Structured Evidential Argumentation System)

² **Note:** At the time of the test SC2 was not a “listed tax avoidance transaction”, it subsequently received attention during Senate hearings focused on KPMG promotions in Oct. 2003, and was recently identified as a listed transaction by IRS: LT 30: [Notice 2004-30](#), 2004-17 I.R.B., Apr 26, 2004 - S Corporation Tax Shelter

2. Could not be classified by looking at tax return data for a single year.
3. Automated scoring systems looking at the affected returns individually were unlikely to flag the issue or identify the associated returns.
4. The entities involved multiple IRS divisions i.e. S-Corp (LMSB), individual beneficiaries (SBSE) and a non-profit entity (TEGE) etc.

Firstly we needed a formal reasoning technique to break down a complex analysis into manageable, systemic subtasks.

Secondly, collaboration with diverse technical experts was required to gain a greater shared understanding of key aspects of the SC2 pattern. Since IRS is geographically dispersed, frequent face-to-face meetings between experts on an on-going basis is highly inefficient, and use of electronic methods such as e-mail and telephone conferences tends to fractionate information exchange and not lead to a group consensus or conclusion. We needed a method for “asynchronous collaboration” to solicit advice, gather evidence, record the developing risk analysis and enable peer review of the work in process.

Thirdly, we needed an efficient way of organizing and sharing information from disparate sources such as tax returns, emails, work papers, technical documents, etc. into a centralized directory.

Fourthly, we needed to classify and record this knowledge in a form that can be rationally compared, searched and referenced at a later date.

SRI International was approached to discuss approaches and technology that are available to meet the Compliance Lab needs. This discussion led to our joint investigation of the use of SEAS software for use in the assessment of compliance risk within the IRS. The Proof of Concept study involved the development and evaluation of exploratory prototypes.

The Compliance Lab developed the concept of **Compliance Risk Patterns (CRiSP)** to describe linked structures and/or transactions that are detectable in IRS filing and compliance databases. CRiSP’s need to be classified and analyzed as a means of detecting high-risk compliance behavior. For example, in the case of the SC2 shelter the CRiSP characteristics include an S-Corp allocating a high proportion of its taxable income to a non-profit entity such as a public employee pension plan. High risk CRiSPs consist of structures and/or transactions similar to those associated with known tax avoidance transactions (such as “listed transactions”).

SRI and the Compliance Lab assembled a multi-disciplinary team of IRS personnel to identify and design the risk analysis components for a CRiSP using the SEAS approach. The team included research and field staff

(Revenue Agents, Technical Subject Matter Experts, International Examiners and Managers) from several IRS divisions (LMSB, SBSE, OTSA – Office of Tax Shelter Analysis, and PFTG – Pre-Filing Technical Guidance). The team analyzed a current compliance issue, built a structured argument prototype to assess the compliance risk for IRS, and evaluated the resulting SEAS prototype.

With the help of the National Research Office, the Compliance Lab setup an IRS SEAS server in Washington, DC for use by the team via the IRS Intranet. Since SEAS is web-based, no additional installation was required. Participants accessed and reviewed the prototype SEAS argument using a standard web browser.

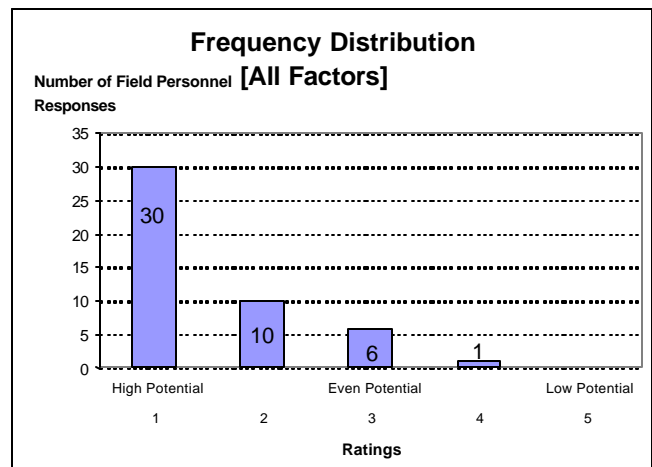
Participants evaluated the SEAS approach based on whether it demonstrated a “high potential payoff” for the IRS for each of the following:

- **Analysis** of compliance risk cases.
- **Understanding** someone else’s analysis.
- **Collaboration** between Field Agents, Technical Advisors, Office of Tax Shelter Analysis & LDC.
- Creating a **corporate memory** for CRiSP cases.
- Encouraging cross-functional learning and **knowledge sharing**.
- **LMSB Issue Management**.

Test Methodology

The test approach emphasized close coordination between contractor, research and field experts in evaluating the technology. This proved to be extremely effective. Careful joint planning of the project using timelines, close supervision by research, and hands on project management by a field manager enabled this test to be completed in four months, on budget, and with judicious use of valuable, IRS technical experts.

Table 1: Field Panel Evaluation



The Results

As shown in Table 1, overall field panel evaluation was highly positive. All participants believed that there is a high potential payoff for IRS use of the SEAS approach. Structured argumentation as a formalized reasoning technique was shown to be effective in improving the efficiency in several areas of the examination process by:

- Breaking down the complexity of the SC² shelter into manageable subtasks.
- Utilizing specialized expertise in a collaborative manner.
- Improving the shared understanding of the compliance issues.
- Detecting and classifying the compliance risk pattern.
- Capturing a point in time analysis that can be recalled and compared at a later stage.

Benefits to IRS

Based on the results of the Proof of Concept study and other templates created in the Compliance Lab, the SEAS approach appears to have good potential to assist IRS to develop enforcement strategies that:

- Systematically assess compliance risk associated with related return groupings.
- Enable collaboration among IRS experts to more rapidly identify and analyze complex schemes.
- Provides access to evidence from multiple sources for multi-disciplinary teams to weigh and agree on an appropriate response.
- Provides auditors with access to more current and comprehensive knowledge about related entities and potential compliance issues that affect the entity that they are assigned to examine.

The professional skill and experience of IRS enforcement employees are most valuable IRS resource. Improved collaboration and accumulation of an accessible “corporate memory” of complex issue & entity risk assessment will improve employee effectiveness and satisfaction

If the IRS can make complex compliance risk assessment more effective and efficient it will help deter corporate and wealthy taxpayers and their tax advisors from the use and promotion of high-risk tax avoidance strategies that exploit complexity and “good optics”.

B: SEAS Technology Overview

The survival of an organization 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 may resort to informal methods, sacrificing structure and rigor. SRI has developed 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.

Introduction to SEAS

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).

SRI International developed the SEAS to encourage the use of structured argumentation which built on the first SEAS prototype (Stokke et al. 1994), – an early warning system for project management.

For this project with the IRS, SEAS was applied to the problem of compliance risk assessment. Our goal was to construct a prototype capable of aiding IRS analysts in leveraging analytic products (SEAS arguments) and methods (SEAS templates) 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 compliance risk assessment was the focus of this prototype, that the tools and methods being developed may have broad application outside of compliance risk assessment as well. 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 e.g. complex risk management.

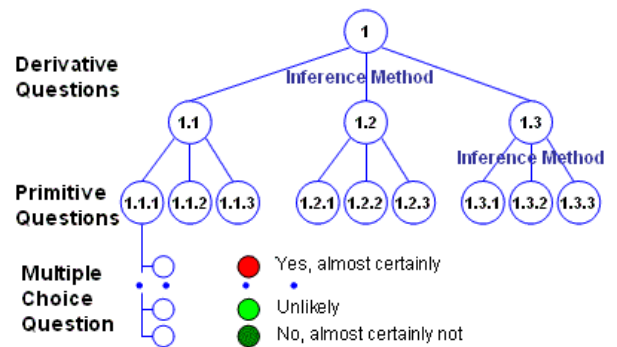
CAPTURING ANALYTIC METHODS

The concept of a *structured argument* is based on a hierarchically organized set of questions (a tree) that is used to assess whether risk of a given type is significant.

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 risk 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 thirteen-question argument template, with nine primitive questions and four derivative questions. Note that question 1 is answered based upon the answers to 1.1, 1.2, and 1.3, and

1.2 is answered based upon the answers to 1.2.1, 1.2.2, and 1.2.3. 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.

Figure 1: An example argument skeleton



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.

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 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.

Figure 2: Argument template question hierarchy

ECONOMIC SUBSTANCE : Does the nature of the entity reflect reasonable business purposes?

1 NORMAL OPERATING BUSINESS : Do the component entities appear to be normal operating businesses?

1.1 UNUSUAL TRANSACTION : Is the evidence of a large, unusual and questionable transaction?

1.2 ECONOMIC REALITY : Does this entity lack economic reality?

1.3 SPECIAL PURPOSE ENTITY : Is this an apparent Special Purpose Entity?

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, and red). Here green might correspond to true, red to false, and the other three to varying degrees of 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.

In general, *discovery tools* are recommended methods for acquiring information relevant to answering questions in an argument template. As illustrated in Figure 4, the evidence could be derived from various sources e.g.

database query, link analysis, parameterized launches of other analytic tools (search engines or specialized applications), or references to other 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.

Finally, a *situation descriptor* describes the type of situations for which the template is intended to be used. Most 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 compliance risk assessment, these could include say, the type of tax entity being analyzed (e.g., a partnership, corporation, trust etc.), the principal actor (e.g., the promoter, the participant etc.), the type of compliance risk pattern (e.g., transitory partnership, basis shifting, leasing related etc.), and the time period (tax years affected). 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.

The high-level templates, as illustrated in Figure 2, derived from the prototype effort are a useful starting point for more in-depth analyses. We imagine that variants of this high-level template will eventually be supported by cascaded templates that are more pointed. In general, while the high-level template is useful in organizing the analysis and reminding analysts of the full range of indicators that need to be assessed, cascaded templates (under the 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, templates capture and deliver best practice.

CAPTURING ANALYTIC PRODUCTS

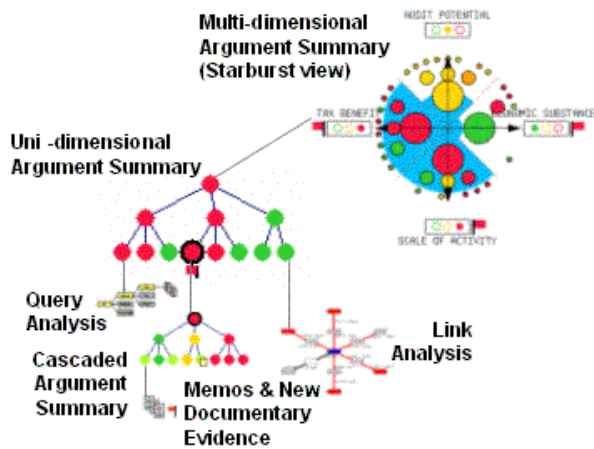
Arguments are formed by answering the questions posed by a template and attaching the evidence that was used in arriving at the selected answers. In essence, an argument organizes the indications and warning signs for the given type of compliance risk.

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

Figure 4: Cascaded structured arguments



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.

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 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. Also when discovery tools are present, they can be used to aid in the collection of evidence.

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.

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.

Figure 5: SEAS argument in browser client

The screenshot displays the SEAS browser client interface for an argument titled "TAX BENEFIT.Generic-CRIsP-test-SC2". At the top, there is a network diagram with red and green nodes and a "No" symbol. Below the diagram is a navigation bar with "print" and "close" buttons. The main content area is divided into sections: "Base Question" with the text "CURRENT YEAR IMPACT: What is the impact on the current year tax rate?", "Analyst" information for Stanley, David J, LMSB Research, dated 27 Feb 2004, and a legend for ownership percentages. The "Rationale" section explains the tax impact of a 90% ownership transfer. Below this are sections for "Documentary Evidence" and "Exhibits", each with a list of items and a "Manual" dropdown menu. The "Discovery Tools" section is currently empty.

A CORPORATE MEMORY OF ANALYTIC KNOWLEDGE

SEAS has been developed as a Web server to communicate 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-based management system, with ephemeral views served up upon demand. Figure 5 shows one such view of a primitive question within an argument.

If we are to recognize emerging issues, then we must relate the present to compliance risk assessments of the past. Additionally, we must understand how the current situation is like or unlike previous situations; how previous assessments evolved and thereby how the present situation might evolve; and how previous analyses were leveraged, mitigated, 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, IRS analysts can leverage the

thinking from the past and present when addressing new cases of compliance risk. 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) [12], CIM (Veridian 2001) [20], 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.

CONCLUSIONS

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. SEAS methodology

- Encourages careful analysis, by reminding the analyst of the full spectrum of indicators to be considered
- Eases complex argument comprehension and communication by allowing multiple visualizations of the data at different levels of abstraction, yet still allowing "drill down" along the component lines of reasoning to discover the details and rationale of others' arguments
- Invites and facilitates argument comparison by framing arguments within common structures
- Allows analysts to readily comprehend the thinking of others by retaining direct links to the source material and its interpretation relative to the conclusions drawn. This memory of analytic thought retains the analytic methods and products of an organization, allowing analysts to leverage the thinking of others both past and present.

We believe that our prototype as implemented in SEAS has shown that the addition of even minimal structure into the analytic process can aid IRS analysts in developing, communicating, explaining, and comparing analytic results.

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 prototype project by SRI, particularly Tom Boyce and to

all IRS staff supporting the effort, in particular, David Stanley for spearheading the effort, Marlon Moore, Francis Cappelletti, Dick Leong, and the IRS senior management team, Don McPartland and Dick Teed for their support.

REFERENCES

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