Argument schemes and provenance to support collaborative intelligence analysis

Alice Toniolo, Federico Cerutti, Nir Oren and Timothy J. Norman
Department of Computing Science, University of Aberdeen,
Scotland, UK
e-mail: {a.toniolo, f.cerutti, n.oren, t.j.norman}@abdn.ac.uk

***Abstract,*—Intelligence analysis is the process of interpreting scattered information to form hypotheses and testing those against evidence. Collaboration enhances this process reducing effort and bias of an individual, while permitting more criticism and different perspectives to be considered. Existing analytical tools support an analyst in collecting information and evaluating hypotheses. However, for effective collaboration, analysts must work together in forming hypotheses from information. We propose an evidential reasoning service that aims at improving collaborative sense-making. It exploits argumentation schemes for structuring annotation and analysis of information, and records provenance to track data and reviews. This service aims to support the analysis by introducing automated reasoning about competing evidence for identifying plausible hypotheses. It provides a uniform reasoning structure that permits integration and facilitates sharing of analyses from different contributors.**

Keywords—intelligence analysis; argument schemes; provenance; collaborative applications.

I. Introduction

Intelligence analysis is an iterative process of interpreting information and elaborating evidence about situations and events [4]. The product of the analysis is a report discussing hypotheses supported by evidence. These reports can be used by decision makers in informing strategies, tactical operations, non-kinetic activities, additional collection requirements, trend analysis, and so on. Collaboration is common among analysts of an intelligence agency [6], however analysts may form coalitions across different organisations in order to address complex tasks. Collaborative analysis is a difficult task within, and more so across, agencies because analysts have access to diverse sources of information that may report conflicting data and their analysis may have different purposes. Further, analysts have different expertise, part of the analysis may be dependent upon other contributors and the analytical approach may differ from agency to agency. A system that supports collaborative annotation and sense-making of information is crucial to deliver timely and accurate intelligence reports.

In order to reduce analyst workload, software tools, which logically organise information, have been developed. TRELLIS [5], for example, focuses on annotating information received from different sources, highlighting contradictions and trustworthiness of sources. XIP-Cohere [2] supports mixed automatic and human annotation. However, very few tools permit collaboration in the analytical process. Entity Workspace [1] supports collaboration in comparing and deciding upon the most likely hypothesis. POLESTAR [10] instead allows an individual portfolio of analysis to be shared across different users that can make suggestions. Although these tools have been positively adopted for collaborative analysis, they focus on *sharing* and *functional* collaboration [6], those being the activities of managing information and expertise and editing reports for completing the analysis. In our research we address collaboration at the *content* level, whereby analysts work together to reason about information and evidence. This type of collaboration permits the elaboration of more hypotheses and enables greater criticism in reasoning. However, problems with integrating and maintaining partial intelligence analyses made by analysts that differ in capabilities, access to information and analytical approaches make this a difficult challenge.

Our core research question is: *How can we support analysts within a coalition to collaborate in sense-making at the content level of analysis?* We propose an evidential reasoning service to enhance content analysis. This service follows the iterative approach of analysts to intelligence analysis and uses *argumentation schemes* for structuring annotation of information and reasoning about competing evidence to build hypotheses. Within argumentation theory, argumentation schemes are used for structuring critical thinking about a controversial statement [3]. Intuitively, these schemes provide structures for making inferences, and for exploring evidence for or against a claim [12]. Assume that an analyst investigating criminal activities asserts that "Jill collaborates with Bob and Bob is a smuggler, then Jill is a smuggler". This inference may be seen as an argument. The fact that Jill is a smuggler may be tentatively accepted unless other evidence dismissing this claim is collected. The structure of this logical inference is extracted and represented as a reasoning pattern that constitutes an argumentation scheme. In addition, the service records *provenance* (i.e., origin of the information) to track data and reviews for assessing the quality of the analytical workflow.

Using argument schemes, our evidential reasoning service guides analysts to uniformly structure links and inferences amongst information. Such structure can be used for autonomous reasoning about competing evidence to assist users in the construction and validation of hypotheses. The service facilitates the integration of partial analyses from different contributors improving collaboration at the content level of analysis. Moreover, it considers provenance that will enhance sharing and maintenance of these analyses. In this paper we discuss: challenges of the reasoning service in Section II; argumentation schemes and provenance for collaborative analysis in Section III; and future work in Section IV.

II. Challenges of the Reasoning Service

Our evidential reasoning service aims at assisting analysts to collaborate throughout the reasoning process in order to create timely and accurate intelligence reports. We envisage our service to be adopted within a platform that enables

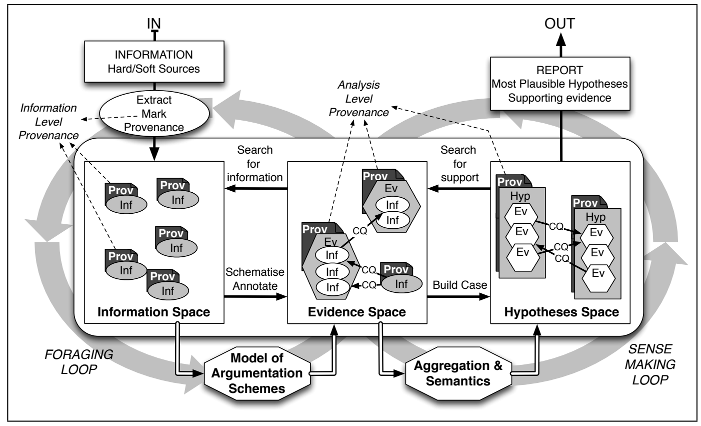

Fig. 1: Overview of the Evidential Reasoning Service

analysts to work together providing data-stream maintenance and managing the contributions to the analysis. The initial information is extracted from soft or hard sources and expressed via a human/machine readable language permitting both user and automatic elaboration (such as in [2]). Although there is no established computational representation of intelligence elements, we can broadly identify *entities* such as people and places, *events* including actions or activities, and *facts* about situations. The analysis focuses on making sense of scattered information by exploring relations between these elements. Important relations are *temporal* relations representing correlation, causality and chronological distribution of events and activities, and *association* relations that connect people and events. The mental process of intelligence analysis is identified by two main loops [11]: the *Foraging Loop* intended as the process of collecting information and extracting evidence; and the *Sense-Making Loop* that aims to structure and annotate evidence to form plausible hypotheses and prepare documents to be presented to the decision makers. In this research we focus on the formation of evidence ("Ev") from information ("Inf") and the formation of hypotheses ("Hyp") from evidence (see Figure 1). Following this model, our service provides a virtual space for each level of elaboration (for a topic) that can be accessed (as appropriate) by a team of contributors and it must proactively assist analysts in these phases.

Within a coalition, experts in different fields have access to different sources, and to a vast amount of domain specific information. The identification of connections is complex because analysts may only be allowed to view or share partial information, they may exploit different evidence, and draw different or even incompatible conclusions. Biases such as confirmation bias, whereby an analyst considers only information that confirms one's beliefs, may also prevent an analyst to draw accurate conclusions. Moreover, teams of analysts may iteratively or concurrently review the analysis and modify hypotheses according to new information. The workflow of partial analyses must be considered to better assess their quality. In the development of the evidential reasoning service we aim to address the following questions:

- How do we support analysts in structuring information and making links by exploring the relations among information?
- How do we support analysts in summarising and annotating evidence for building hypotheses?
- How do we support analysts in identifying acceptable hypotheses?
- How does the history of analysis and manipulation affect the acceptability of claims for an analyst?

III. Argumentation Schemes and Provenance

Our approach to the above questions is to employ argumentation schemes for structuring the analysis and provenance for recording the workflow of the analysis.

Argumentation schemes. An argumentation scheme is a structured way of making presumptive inferences, stating explicitly what the premises are and what conclusions can be drawn from these premises. Associated with an argumentation scheme are critical questions (CQs), which can be used to challenge the validity of arguments. Argumentation schemes represent a method for formulating arguments in argumentation theory, whose computational aspects revealed to be of great interest since the seminal work of [3]. Argumentation theory has increasingly received attention in artificial intelligence as a mechanism for representing autonomous reasoning with uncertain and incomplete information [8], by providing methods for deriving the acceptability status of arguments. In [3] an argument is rationally acceptable if it is defended against attacking arguments. For example, consider argument A_1, "Jill collaborates with Bob, Bob is a smuggler, thus Jill is a smuggler too". An attack is an argument A_2, "Jill had no contact with Bob, thus Jill does not collaborate Bob". The fact that Jill is a smuggler cannot be rationally accepted since A_2 attacks A_1. However, if A_2 is attacked by a new argument A_3, "Jill and Bob has been introduced by Mark, therefore Jill and Bob had a direct contact", claim A_1, defended by A_3, may then be reinstated. In an argumentation framework, where arguments and attack and support relationships between arguments are specified, a criterion, called semantics, is defined to establish the acceptability of arguments considering relationships between arguments.

In computational systems the use of argumentation schemes has been introduced to formulate arguments since they provide structures that can be applied to diverse information and permit the representation of reasoning in complex domains [8]. Empirical studies in domains such as law provide a variety of argumentation schemes derived from patterns of common human reasoning and dialogue [12]. Argumentation schemes are defeasible in the sense that the premises warranting the conclusion are tentatively accepted according to existing evidence. In the light of new received information the conclusions may be discarded if this invalidates the claim. In our example, argument A_1 can be structured using the scheme (from [12]):

- **Premise 1** - M member of group G has quality Q,
- **Premise 2** - if M has property Q, G will have Q,

$\Rightarrow$ **Conclusion** - Therefore, every member of G have Q.

where quality Q is being a smuggler, M is Bob and the group G is formed by Bob and Jill. This scheme can be applied to form a different argument with similar structure for example "A batch of goods M in the warehouse G is smuggled, thus all the goods in the warehouse G are smuggled".

We believe that argumentation schemes can be used in collaborative intelligence analysis for structuring information to form evidence in terms of relations between entities, events and facts (highlighted in Figure 1). The service should assist users in selecting schemes that may fit the information available. Furthermore, we will investigate methods to aggregate and combine schemes to summarise interpretations of evidence for assisting analysts in building hypotheses. In this process critical questions play an important role. CQs may identify supporting relations among information or may highlight missing information such as premises that are necessary for a specific inference to be supported, but for which there is insufficient evidence. CQs make explicit links to alternative explanations, and may prevent confirmation biases by highlighting evidence against the most favourable hypotheses, ensuring a more objective evaluation. Similar to argumentation schemes, other structures have already been tested with existing tools for intelligence analysis. For example, TRELLIS [5] presents predefined constructs for asserting statements and associated reasons. However, argumentation schemes provide structure to those reasons that can be analysed by automated reasoning mechanisms to assist analysts in drawing conclusions. The defeasible of arguments is suitable to handle the dynamic intelligence analysis process, where new information is continuously being collected. An argumentation semantics can be used to assist analysts in deciding the status of the hypotheses.

Provenance. For effective collaborative analysis, analysts examining different information must consider *how, when, where* this information has been gathered and *by whom* it has been manipulated. Existing toolkits provide support for automatic or manual provenance annotation following data-model recommendations (e.g., PROV-DM [7]), and provide repositories for storing such metadata and querying services to extract relevant information [9]. We will employ such models to annotate the provenance of incoming data and the process of analysis. In particular, as shown in Figure 1, we will record provenance at different levels: *a)* the information level, where provenance of documents and information is stored, including sources and context of collection; and *b)* the analysis level, where provenance is a record of each phase of the analysis for a topic, including in particular sources and contributors, the creation of new schemes, data used, timestamps and updates.

This contextual information is used to retrieve the history of analysis and it must be integrated with the reasoning service in order for an analyst to better assess claims. In fact, understanding the workflow of analysis may affect conclusions and may lead to discard some hypotheses or evidence because they come from flawed reasoning processes. Argumentation schemes will be designed for relating inferences to, for example, the expertise of the source, when the claim was made, the temporal consistency of claims, expert reliability, trustworthiness of a source, and so on. Such schemes can be used as bridge between the reasoning process and the analysis of provenance records. To date, however, provenance within argumentation is yet to be explored. Furthermore, provenance has important applications in maintaining large volumes of heterogeneous data, but it has not been integrated within intelligence analysis tools. An introductory work is proposed in [13] for visual analytics. Although this work proposes a layered method to record analysis, here, we focus on the introduction of provenance into the reasoning process, not only as an external record that an analyst may consult. The analysis of contextual information contributes to the overall sense-making process in collaborative analysis.

Example. We introduce, here, an example to illustrate relations amongst information, how a scheme will support the generation of hypotheses, and how the introduction of contextual information affects the analysis. The goal is to establish the presence of any criminal activities across the border in a named area of interest, between locations L1 and L2 (L1-L2). The first scheme discussed is an *abductive argument from effects to cause* [12], which explores the causal relation between a set of facts F, and its plausible causal explanation C. Statements C and F may be collected from different sources, added by the analyst or suggested by the system. Here, the scheme is used by an expert of video recordings to state that some aerial images show that there is a gang G suspected of smuggling forbidden products P1 across the border L1-L2, and this is a possible explanation of why product P1 arrived in L2.

- **Premise 1** - The set of events $F = \{$"A: there is a forbidden product P1 in location L2", "B: a gang G smuggles products P1 in location L1"$\}$ has been recorded,
- **Premise 2** - "C: Smugglers G crossed the border L1-L2" is the best satisfactory causal explanation of F so far,
- ⇒ **Conclusion** - Therefore, C is plausible as the cause of F.

This scheme can be challenged with the question "How strong is the explanation C?" or "Is there any better explanation?". This partial analysis is shared with an analyst specialising in trafficking within the region. The new scheme is an *abductive argument from evidence to a hypothesis* [12]. Here, the hypothesis is that, since the smugglers crossed L1-L2, criminal activities have occurred at the border.

- **Premise 1** - If hypothesis "H: There are criminal activities at the border checkpoint L1-L2" is true then "C: Smugglers G crossed the border L1-L2" will be observed to be true,
- **Premise 2** - C has been observed to be true,
- ⇒ **Conclusion** - Therefore, hypothesis H is true.

A question that links to the previous analysis is "Has C been observed to be true?". Other questions may exploit alternative reasons for C being true; e.g., that the smugglers bypassed the checkpoint, thus the suspected activities are not associated with the checkpoint. Moreover, introducing information from provenance records about the reliability of the sources of F can invalidate the hypothesis if not supported by stronger evidence. The *argumentation scheme from expert opinion* may be used to relate a statement to its source [12]. Assume that the fact that product P1 was found in L2 has been asserted by E, an expert in identifying illegal products.

- **Premise 1** - Source E is an expert in domain "identification of products P1,
- **Premise 2** - E asserts that proposition "A: there is a forbidden product P1 in location L2" is true,
- ⇒ **Conclusion** - Therefore, A may be taken to be true.

In contrast with the others, this scheme permits challenges against both claim A and the association of A with source E. If there is evidence that leads the analyst to think that E is unreliable (e.g., through previous interactions), the claim can be rendered invalid. In fact, the new evidence attacks A and challenges the inferences C and H made on the basis of A.

IV. Future work & Conclusion

Effective collaborative intelligence analysis must be supported at the content level, in reasoning about information and formulating consistent hypotheses. We proposed an evidential reasoning service that guides analysts in structuring the process of analysis employing argumentation schemes, where partial elaborations can be shared more easily and an autonomous reasoning mechanism can be applied to support the sense-making. Recording provenance data permits one to assess the quality of information and analyses. The first step to develop this service is to identify an appropriate set of argumentation schemes and related critical questions to use for the analysis. These will be based on a computational language and a collection of ground facts for representing intelligence. We use existing schemes as a starting point [12], however such schemes may not be sufficiently expressive. We will design new schemes for representing relations amongst information by employing diagramming methods [12] for extracting schemes from corpora of intelligence reports. More importantly we aim to engage with experts in the field for validating and defining relevant schemes. A further question is how to include provenance data within the schemes, such as the temporal persistence of information, the workflow of analysis, the trustworthiness of data, and so on. We will identify argumentation semantics able to reflect how analysts consider some hypotheses acceptable, based for example on their expertise or on the reliability of the supporting evidence in order to ensure an objective evaluation.

We will explore methods for dealing with uncertain information that may come from untrusted sources or sources that have reported unreliable data in the past. We will study how this affects the construction of evidence and the acceptability of hypotheses. Furthermore, the information accessible by an analyst may be incomplete because many reports have limited distribution, clearance levels may dictate what information is available, and so on. Critical questions may help in identifying gaps and prompt the user to seek more information. Provenance may also help by disclosing some contextual information to ensure that if only limited analysis is shared an analyst may at least assess its reliability. However, these options may not always be allowed. Thus, we must study methods for addressing information gaps when sharing is limited. An important issue to be addressed is how users perceive the employment of argumentation schemes. Most of the work on argumentation is for autonomous reasoning, although such schemes have been used to analyse real world arguments [12]. Here, however, argumentation schemes are used in a system that interacts with users to assemble scattered information, and many issues may arise. For example, the support is not adequate for the inferences that the analyst intends to report. In deploying the reasoning service, we must consider what schemes should be suggested to the users, and which ones would best fit the information for providing valid support. We previously claim that our evidential reasoning service aims at improving the sense-making in collaboration at the content level, in order to assess such a claim we must perform experiments in order to test how effective our system is in identifying plausible explanations in comparison to the use of other tools (e.g., [1,10]) according to both the quality of the hypotheses identified and the ease of use of the system.

In conclusion, we believe that our evidential reasoning service based upon argumentation schemes and provenance is suitable to support collaborative intelligence analysis. This service will facilitate the integration of analyses from different contributors permitting a more objective evaluation of hypotheses and reducing the workload of analysts. Furthermore, tracking the use of data within the analytical process may help analysts in understanding the utility of information for planning more focussed collection activities.

Acknowledgment

The authors would like to thank Susan Toth at US Army Research Laboratory for the valuable feedback. Research was sponsored by US Army Research laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

[1] E. Bier, S. Card, and J. Bodnar, "Entity-based collaboration tools for intelligence analysis," Proceedings of the IEEE Symposium on Visual Analytics Science and Technology, 2008, pp. 99–106.

[2] A. De Liddo, Á. Sándor, and S. Buckingham-Shum, "Contested collective intelligence: Rationale, technologies, and a human-machine annotation study," Computer Supported Cooperative Work, vol. 21, 2012, pp. 417–448.

[3] P. M. Dung, "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games," Artificial Intelligence, vol. 77, 1995, pp. 321–357.

[4] B. Fischhoff and C. Chauvin, Intelligence Analysis: Behavioral and Social Scientific Foundations. The National Academies Press, 2011.

[5] Y. Gil and V. Ratnakar, "TRELLIS: An interactive tool for capturing information analysis and decision making," Knowledge Engineering and Knowledge Management: Ontologies and the Semantic Web, vol. 2473, Springer Berlin Heidelberg, 2002, pp. 37–42.

[6] Y.-a. Kang and J. Stasko, "Characterizing the intelligence analysis process: Informing visual analytics design through a longitudinal field study," Proceedings of the IEEE Conference on Visual Analytics Science and Technology, 2011, pp. 21–30.

[7] L. Moreau and P. Missier, "PROV-DM: The PROV Data Model," 2013, Available at: http://www.w3.org/TR/prov-dm/ [retrieved: May, 2013].

[8] N. Oren, T. J. Norman, and A. Preece, "Subjective logic and arguing with evidence," Artificial Intelligence, vol. 171, 2007, pp. 838–854.

[9] E. Pignotti, P. Edwards, A. Preece, N. Gotts, and G. Polhill, "Enhancing workflow with a semantic description of scientific intent," The Semantic Web: Research and Applications, vol. 5021, Springer Berlin Heidelberg, 2008, pp. 644–658.

[10] N. Pioch and J. Everett, "POLESTAR: collaborative knowledge management and sensemaking tools for intelligence analysts," Proceedings of the 15th International Conference on Information and Knowledge Management, 2006, pp. 513–521.

[11] P. Pirolli and S. Card, "The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis," Proceedings of the International Conference on Intelligence Analysis, 2005.

[12] D. Walton, C. Reed, and F. Macagno, Argumentation schemes. Cambridge University Press, 2008.

[13] B. Wong, K. Xu, and S. Attfield, "Provenance for intelligence analysis using visual analytics," Proceedings of the CHI Workshop on Analytic Provenance, 2011.