

CISpaces.org: from Fact Extraction to Report generation¹

Federico CERUTTI^a, Timothy J. NORMAN^b, Alice TONIOLO^c, and
Stuart E. MIDDLETON^b

^a*Cardiff University, UK*

^b*University of Southampton, UK*

^c*University of St Andrews, UK*

Abstract. We introduce CISpaces.org, a tool to support situational understanding in intelligence analysis that complements but not replaces human expertise. The system combines natural language processing, argumentation-based reasoning, and natural language generation to produce intelligence reports from social media data, and to record the process of forming hypotheses from relationships among information. In this paper, we show how CISpaces.org meets the desirable requirements elicited from senior professionals, and demonstrate its usage and capabilities to support analysts in delivering effective and tailored intelligence to decision makers.

Keywords. argumentation, open source intelligence, report generation

1. Introduction

In this paper we introduce CISpaces.org, Collaborative Intelligence Spaces Online, a novel application of argumentation theory to provide analysis capabilities from open sources (OSINT). CISpaces.org is a suite of tools and algorithms for the support of sense-making in intelligence analysis, for the extraction of facts from evidence from social media sources, and for the dissemination of natural language reports. CISpaces.org supports the process of sensemaking, complementing human expertise in the generation of intelligence products. This web-based tool builds on top of argumentation-based systems, combining a structured graphical representation of the analysts reasoning process with efficient algorithms for the automated identification of plausible hypotheses. This tool is integrated with a fact extraction engine, through which open sources can be queried, and information extracted can be brought into an analysis.

¹Submitted to IAT

This project was awarded funding through the Defence and Security Accelerator themed competition Revolutionise the human information relationship for defence under Contract DSTLX1000113927. This research was partially sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. 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 U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

The main contribution of this paper is to propose a system that supports the full data-to-decision process from information extraction, to hypotheses formation, to report generation that can be used for briefings to inform decision-makers. In respect to argumentation systems, CISpaces.org is based on an existing evidential reasoning service [23,22], and it improves upon this with the following novel features: (1) fully web-based; (2) analytic capabilities for extracting facts from Twitter streaming, building on top of [13,14]; (3) automatic reports generation summarising argumentation-based analyses.

CISpaces.org facilitates sensemaking within the intelligence analysis process in a declarative format. Intelligence analysis is an iterative process of foraging for information and sensemaking in which the analysis structure increases incrementally from a shoebox of information, through evidence files, to the generation and evaluation of hypotheses. Differently from existing tools [9,8], CISpaces.org provides a method to record and support the process of forming hypotheses from the relationships among information which enables the analyst to highlight information or assumptions that may lead to interrelated as well as alternative hypotheses. CISpaces.org makes this core process of reasoning explicit, providing further support for structuring reasoning and mitigating biases. The reasoning mechanism identifies what evidence and claims together constitute a plausible interpretation of an analysis.

A running example, introduced in Section 2 will support the presentation of our approach that is based on: argumentation schemes, as tools for supporting critical thinking; ASPIC+; abstract argumentation; and fact extraction from OSINT (Section 3). In Section 4 we show how CISpaces.org satisfies the three elicited desiderata, namely having the ability to deal with open source intelligence; having the ability to manipulate graphically argumentative structures and to evaluate them efficiently; and having the ability to automatically generate a natural language report of the performed analysis. In Section 5, we discuss different venues in which CISpaces.org has generated considerable interest.

2. Motivation and Desiderata

We illustrate the use of CISpaces.org based on an extract of a UK DSTL Non-Combatant Evacuation Operation Scenario [24]. The example we develop in this paper is our interpretation of such scenario concerned with the breakout of a dangerous virus causing the rising of unrest in a fictional city, Squirrel City (Figure 1).

The analysis is concerned with an operation conducted to identify a safe evacuation route for UK Nationals who are in danger in Squirrel City. Assume that analysts have already identified which UK nationals are at risk and where they are, notably around the Stock Market area. There are two main routes for evacuation: one to a port, and one to an airport. To understand in near real-time the situation, the analyst monitors social media,² from which there are anecdotal pieces of evidence indicating that there is a virus outbreak in Squirrel City, that the population is concerned that no actions have been taken, and that groups of people are rioting downtown.

²The micro-blogging posts — around 400 — we used in the creation of this example, come from Twitter posts crawled during the Hurricane Sandy event in October 2012 around New York, re-purposed using text search&replace to look like a virus outbreak or riots rather than a flooding event. This means the posts are “real” in the sense that there are many examples of real examples of bad grammar, poorly structured text and contradictory reports (e.g. New York Stock Exchange flooding reports which are later debunked).




Figure 1. Fictional maps of Squirrel City — based on the fictional city of Solaria (<http://wiki.opengeofiction.net/wiki/index.php/Solaria>, on 9th April 2018) — with highlighted the Stock market, the Port, and the Airport.

As the situation proceeds, analysts need a rigorous means to gather new relevant pieces of information, record and support the process of forming hypotheses from relationships among information, and to communicate hypotheses to decision makers in appropriate tailored reports as required, or as the situational understanding shifts.

To successfully support the analyst in their situational understanding, and building on our previous work showing the effectiveness of using argumentation theory for supporting intelligence analysis, we need to empower them in:

- D1: Extracting facts from social media (e.g. Twitter) and use them in their argumentation-based analysis;
- D2: Manipulating graphically argumentation structures and accessing an efficient algorithm for their evaluation;
- D3: Automatically generating reports from the analysis.

In the following we show how CISpaces.org satisfies those three desiderata, that build on top of existing works in argumentation theory reviewed in the following section.

3. Background

3.1. Argumentation Schemes

Argumentation schemes [25] are abstract reasoning patterns derived from empirical studies of human argument and debate, and further adapted in this work from literature and experts in intelligence analysis [23]. Each scheme has a set of critical questions that represents standard ways of critically probing into an argument to find aspects of it that are open to criticism. For instance, the following is the argumentation scheme for argument from cause *C* to effect *E*:

Major Premise: Generally, if *C* occurs, then *E* might occur.

Minor Premise: In this case *C* might have occurred.

Conclusion: Therefore, in this case *E* might have occurred.

Critical questions are:

CQ1: Is there evidence for *C* to occur?

CQ2: Is there a general rule for *C* causing *E*?

CQ3: Is the relationship between C and E causal?

CQ4: Is there any exception to the causal rule that prevent E from occurring?

CQ5: Has C happened before E?

CQ6: Is there any other cause C' that might have caused E?

The purpose of schemes in CISpaces.org is to guide analysts in drawing inferences, and critical questions are available to analysts as a means to reflect on potential issues during the formation of hypotheses. Instantiated schemes can be mapped to the overall ASPIC framework following the approach proposed in [16].

3.2. ASPIC+ [15]

Assume a logical language \mathcal{L} , and a set of *strict* or *defeasible* inference rules — resp. $\varphi_1, \dots, \varphi_n \longrightarrow \varphi$ and $\varphi_1, \dots, \varphi_n \Longrightarrow \varphi$. A strict rule inference always holds — i.e. if the *antecedents* $\varphi_1, \dots, \varphi_n$ hold, the *consequent* φ holds as well — while a defeasible inference “usually” holds. Arguments are constructed w.r.t. a knowledge base with two types of formulae.

Definition 1. An *argumentation system* is as tuple $AS = \langle \mathcal{L}, \mathcal{R}, \bar{\cdot}, \nu \rangle$ where:

- $\bar{\cdot} : \mathcal{L} \mapsto 2^{\mathcal{L}}$ is a contrariness function s.t. if $\varphi \in \bar{\psi}$ and:
 - * $\psi \notin \bar{\varphi}$, then φ is a *contrary* of ψ ;
 - * $\psi \in \bar{\varphi}$, then φ is a *contradictory* of ψ ($\varphi = \neg\psi$);
- $\mathcal{R} = \mathcal{R}_d \cup \mathcal{R}_s$ is a set of strict (\mathcal{R}_s) and defeasible (\mathcal{R}_d) inference rules such that $\mathcal{R}_d \cap \mathcal{R}_s = \emptyset$;
- $\nu : \mathcal{R}_d \mapsto \mathcal{L}$, is a partial function.

A *knowledge base* in an AS is a set $\mathcal{K}_n \cup \mathcal{K}_p = \mathcal{K} \subseteq \mathcal{L}$; $\{\mathcal{K}_n, \mathcal{K}_p\}$ is a partition of \mathcal{K} ; \mathcal{K}_n contains *axioms* that cannot be attacked; \mathcal{K}_p contains *ordinary premises* that can be attacked.

An *argumentation theory* is a pair $AT = \langle AS, \mathcal{K} \rangle$.

Let us now recall the notion of *argument* in ASPIC+.

Definition 2. An *argument* \mathbf{a} on the basis of a $AT = \langle AS, \mathcal{K} \rangle$, $AS = \langle \mathcal{L}, \mathcal{R}, \bar{\cdot}, \nu \rangle$ is:

1. φ if $\varphi \in \mathcal{K}$ with: $\text{Prem}(\mathbf{a}) = \{\varphi\}$; $\text{Conc}(\mathbf{a}) = \varphi$; $\text{Sub}(\mathbf{a}) = \{\varphi\}$; $\text{Rules}(\mathbf{a}) = \text{DefRules}(\mathbf{a}) = \emptyset$; $\text{TopRule}(\mathbf{a}) = \text{undefined}$.
2. $\mathbf{a}_1, \dots, \mathbf{a}_n \longrightarrow / \Longrightarrow \psi$ if $\mathbf{a}_1, \dots, \mathbf{a}_n$, with $n \geq 0$, are arguments such that there exists a strict/defeasible rule $r = \text{Conc}(\mathbf{a}_1), \dots, \text{Conc}(\mathbf{a}_n) \longrightarrow / \Longrightarrow \psi \in \mathcal{R}_s / \mathcal{R}_d$. $\text{Prem}(\mathbf{a}) = \bigcup_{i=1}^n \text{Prem}(\mathbf{a}_i)$; $\text{Conc}(\mathbf{a}) = \psi$; $\text{Sub}(\mathbf{a}) = \bigcup_{i=1}^n \text{Sub}(\mathbf{a}_i) \cup \{\mathbf{a}\}$; $\text{Rules}(\mathbf{a}) = \bigcup_{i=1}^n \text{Rules}(\mathbf{a}_i) \cup \{r\}$; $\text{DefRules}(\mathbf{a}) = \{d \mid d \in \text{Rules}(\mathbf{a}) \cap \mathcal{R}_d\}$; $\text{TopRule}(\mathbf{a}) = r$

An argument can be attacked in its premises or its conclusion.³

Definition 3. Given \mathbf{a} and \mathbf{b} arguments, \mathbf{a} *defeats* \mathbf{b} iff \mathbf{a} *successfully rebuts* or *successfully undermines* \mathbf{b} , where:

- \mathbf{a} *successfully rebuts* \mathbf{b} (on \mathbf{b}') iff $\text{Conc}(\mathbf{a}) \notin \bar{\varphi}$ for some $\mathbf{b}' \in \text{Sub}(\mathbf{b})$ of the form $\mathbf{b}'_1, \dots, \mathbf{b}'_n \Longrightarrow \neg\varphi$;
- \mathbf{a} *successfully undermines* \mathbf{b} (on φ) iff $\text{Conc}(\mathbf{a}) \notin \bar{\varphi}$, and $\varphi \in \text{Prem}(\mathbf{b}) \cap \mathcal{K}_p$.

³In the following we will not make use of preference and of undercuts on the inference rules.

Arguments and attacks together form an abstract argumentation framework.

Definition 4. *AF is the abstract argumentation framework defined by $AT = \langle AS, \mathcal{K} \rangle$, $AS = \langle \mathcal{L}, \mathcal{R}, \bar{\cdot}, \nu \rangle$ if \mathcal{A} is the smallest set of all finite arguments constructed from \mathcal{K} satisfying Def. 2; and \rightarrow is the defeat relation on \mathcal{A} as defined in Def. 3.*

3.3. Abstract Argumentation Framework

An argumentation framework [4] consists of a set of arguments and a binary attack relation between them.

Definition 5. *An argumentation framework (AF) is a pair $\Delta = \langle \mathcal{A}, \rightarrow \rangle$ where \mathcal{A} is a set of arguments and $\rightarrow \subseteq \mathcal{A} \times \mathcal{A}$. We denote with $\mathbf{a}_2 \rightarrow \mathbf{a}_1$ when $\langle \mathbf{a}_2, \mathbf{a}_1 \rangle \in \rightarrow$.*

The basic properties of conflict-freeness, acceptability, and admissibility of a set of arguments are fundamental for the definition of argumentation semantics.

Definition 6. Given an AF $\Delta = \langle \mathcal{A}, \rightarrow \rangle$:

- a set $S \subseteq \mathcal{A}$ is a *conflict-free* set of Δ if $\nexists \mathbf{a}_1, \mathbf{a}_2 \in S$ s.t. $\mathbf{a}_1 \rightarrow \mathbf{a}_2$;
- an argument $\mathbf{a}_1 \in \mathcal{A}$ is *acceptable* with respect to a set $S \subseteq \mathcal{A}$ of Δ if $\forall \mathbf{a}_2 \in \mathcal{A}$ s.t. $\mathbf{a}_2 \rightarrow \mathbf{a}_1$, $\exists \mathbf{a}_3 \in S$ s.t. $\mathbf{a}_3 \rightarrow \mathbf{a}_2$;
- a set $S \subseteq \mathcal{A}$ is an *admissible* set of Δ if S is a conflict-free set of Δ and every element of S is acceptable with respect to S .

An argumentation semantics σ prescribes for any AF Δ a set of *extensions*, denoted as $\mathcal{E}_\Delta(\sigma)$, namely a set of sets of arguments satisfying the conditions dictated by σ . Here we need to recall only definition of preferred (denoted as \mathcal{PR}) semantics only.

Definition 7. Given an AF $\Delta = \langle \mathcal{A}, \rightarrow \rangle$, a set $S \subseteq \mathcal{A}$ is a *preferred extension* of Δ , i.e. $S \in \mathcal{E}_\Delta(\mathcal{PR})$, iff S is a maximal (w.r.t. set inclusion) admissible set of Δ .

An argument \mathbf{a}_1 is *credulously* accepted with respect to a given semantics σ and a given AF Δ iff \mathbf{a}_1 belongs to at least one extension of Δ under σ : $\exists E \in \mathcal{E}_\Delta(\sigma)$ s.t. $\mathbf{a}_1 \in E$. An argument \mathbf{a}_1 is *skeptically* accepted with respect to a given semantics σ and a given AF Δ iff \mathbf{a}_1 belongs to each extension of Δ under σ : $\forall E \in \mathcal{E}_\Delta(\sigma)$ $\mathbf{a}_1 \in E$. Let us denote with $\mathcal{E}_\Delta^{\mathcal{S}}(\sigma) = \{\mathbf{a} \in \mathcal{A} \mid \mathbf{a} \in E, \forall E \in \mathcal{E}_\Delta(\sigma)\}$.

Let us now define a function that given a set of arguments returns their conclusions, i.e. *Conclusions* : $2^{\mathcal{A}} \rightarrow 2^{\mathcal{L}}$, such that given a set of arguments $E \subseteq \mathcal{A}$, $\text{Conclusions}(E) = \{\text{Conc}(\mathbf{a}) \in \mathcal{L} \mid \mathbf{a} \in E\}$. In particular, $\text{Conclusions}(\mathcal{E}_\Delta^{\mathcal{S}}(\sigma))$ is the set of conclusions of skeptically accepted arguments, which — as we demonstrate in Section 4.3 — is not necessarily equivalent to the intersection of credulously accepted conclusions, i.e. $\{\text{Conc}(\mathbf{a}) \in \mathcal{L} \mid \exists E \in \mathcal{E}_\Delta(\sigma), \mathbf{a} \in E\}$.

Preferred extensions are computed as per [1, Algorithm 1], that is at the basis of the newest version of the ArgSemSAT algorithm that won the first prize in the preferred semantics tracks at the second International Competition on Computational Models of Argumentation (ICCMA 2017).

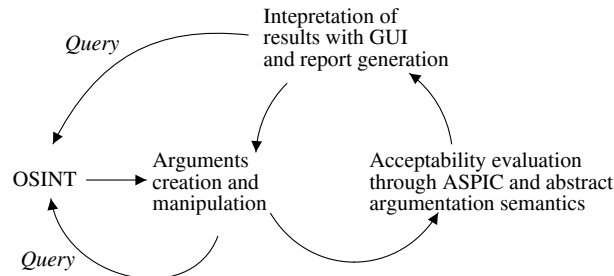


Figure 2. Overview of CISpaces.org functionalities.

3.4. Fact Extraction from Social Media in the Context of Argumentation Mining

Argumentation mining [7] has focused on the extraction from textual documents sets of claims, supporting premises and optionally any associated conclusions or propositions. It has been applied to social media, on platforms such as Twitter. For example [11] experimented with classifying tweets into argumentative categories for claim, claim with evidence, counter claim, counter claim with evidence. The dataset used was the annotated London riots twitter dataset [17]. They used a bag of features made up of unigrams, punctuation, and parts-of-speech (POS) tags, and a Support Vector Machine classifier, and reported a precision of 0.86 which is probably too high since the dataset contained about 60% retweets.

Social media is challenging from a rhetoric point of view, with its informal structure, variable syntax and common spelling errors. In this paper, we will be focusing on extracting the actual propositional structures (entities and relations) associated with statements that might contain supporting or debunking evidence for claims. We are not trying to classify social media posts as containing a claim and evidence. Posts with possible evidence are automatically aggregated, statistically grouped based on entities and relation occurrence frequency, and presented to the user for their eventual inclusion in the analysis.

4. CISpaces.org and Desiderata

CISpaces.org is a web-based system that enables an analyst to analyse OSINT by: (1) graphically manipulating arguments based upon pieces of information gathered from OSINT sources as well as analysts claims; (2) evaluating the arguments via ASPIC and abstract argumentation semantics; (3) reviewing the acceptability of claims and pieces of information via the graphical user interface and an automatically generated report. This is summarised in Figure 2.

We now review how CISpaces.org satisfies the desiderata highlighted in Section 2.

4.1. Desideratum 1: Open Source Intelligence

To provide a proof-of-concept for extracting data from social media, we have integrated into CISpaces.org an OSINT crawler and analytics engine to perform near real-time fact extraction. Since it is beyond the scope of this paper to assess the quality of the fact

extraction system we implemented, we refer an interested reader to the relevant literature cited below. Our focus is, indeed, to demonstrate the possibility of extracting factual claims within social media sources such as 'BBC report 20 dead in New York subway', which can then be analysed by the human intelligence analyst and if relevant introduced as evidence to support or reject hypotheses within an argumentative analysis.

The fact extraction algorithm is based on classic previous work in Open Information Extraction (OpenIE) such as OLLIE [12], WOE [26] and OpenIE5 [3]. We have used experience gained from previous information extraction work [13,14] using syntactic and linguistic pattern analysis to tailor our approach. The bootstrap algorithm is based on parts of speech (POS) patterns first used by ReVerb [6], restricted to proper nouns only so they produce a high precision, but low recall, set of seeds for an unlabelled corpus of training posts. The noun and rel values in each seed are then used as terminal points for a linguistic dependency graph walk of each post to generate an open extraction template that can be executed on future unseen posts. We use the Stanford Dependency Parser to get our dependency trees. We performed a linguistic analysis of relevant OSINT posts typical for military scenarios and used the findings to optimize our dependency graph walk algorithm.

Once the open extraction templates are created the fact extraction algorithm is ready to perform real-time extractions. Each extraction template generates a set of (noun, rel, noun) type extractions, which are cross-checked with the keywords of interest as specified by the human intelligence analyst. This allows us to generate in near real-time factual claims which are then aggregated and presented to the CISpaces.org interface as OSINT information clustered by mention frequency. Each factual claim keeps a provenance link back to the original post, so the human intelligence analyst can see the original context and source of the information.

Currently, CISpaces.org allows an analyst to query and import facts extracted from OSINT as atomic statements. In future developments, we envisage the system to be able to import partial analyses in the form of argument graphs, which may for example be extracted via argument mining tools [5] or gathered through conversational agents [22].

Facts that have been extracted that are of relevance for our running example are *Unseen peek of A&E admissions in downtown Medical Centres*, *Large groups of people are leaving downtown*, *Media report virus outbreak in Squirrel City*, *Media report that population is concerned that no actions are taken to mitigate threats*, and *Local population rioting downtown*.

4.2. *Desideratum 2: Graphical Manipulation of Arguments*

The CISpaces.org interface enables a user to draw a graph of arguments, formed by premises and conclusions, and linked via inference or conflict nodes. The result is a directed graph, referred to as *WDG*. Based on the AIF format [2], CISpaces.org uses two distinct types of nodes: information nodes (or I-nodes) and scheme nodes (or S-nodes). Information nodes are further distinguished in two sub-sets, Facts and Claims, which are useful for analysts to broadly distinguish between information derived from external sources (facts) and observations drawn by analysts during the analysis (claims). S-nodes, as per AIF, are divided in two sub-sets: rule of inference application RA-nodes and conflict application CA-nodes, respectively represented as Pro and Con nodes.

In the graph of arguments *WDG* produced by analysts within the CISpaces.org interface, nodes are connected by edges whose semantics is implicitly defined by their use.

For instance, an information node connected to a Pro scheme node, with the arrow terminating in the latter, would suggest that the information node serves as a premise for an inference rule. The graph $WDG = \langle N, \rightsquigarrow \rangle$, is formed by N , the set of nodes, and \rightsquigarrow , the set of edges. In order to map WDG to an ASPIC system, we assume the following:

- $P \subseteq N$ is the set of I-nodes, where each I-node in the graph is written p_i ;
- An S-node is referred to as ℓ_{type} where $type = \{Pro, Con\}$;
- An inference rule is $[p_1, \dots, p_q \rightsquigarrow \ell_{Pro} \rightsquigarrow p_\phi]$ indicating that p_1, \dots, p_q are incoming nodes and an outgoing node p_ϕ
- Conflict schemes can be either $[p_1 \rightsquigarrow \ell_{Con} \vdash p_2]$ or $[p_1, \dots, p_n \rightsquigarrow \ell_{Con} \rightsquigarrow p_\phi]$.

Given a $WDG = \langle N, \rightsquigarrow \rangle$, its corresponding ASPIC+ system $AS = \langle \mathcal{L}, \mathcal{R}, \bar{\cdot}, \nu \rangle$ is:

- $\forall p \in P \subseteq N, p \in \mathcal{L}$;
- $\mathcal{R} = \mathcal{R}_s \cup \mathcal{R}_d$ with $\mathcal{R}_s = \emptyset$ and $\forall [p_1, \dots, p_n \rightsquigarrow \ell_{Pro} \rightsquigarrow p_\phi], p_1, \dots, p_n \Rightarrow p_\phi \in \mathcal{R}_d$;
- $\forall [p_1 \rightsquigarrow \ell_{Con} \rightsquigarrow p_2], p_1 \in \overline{p_2}$;
- $\forall [p_1, \dots, p_n \rightsquigarrow \ell_{Con} \rightsquigarrow p_\phi]$, is mapped as $p_1, \dots, p_n \Rightarrow p_h \in \mathcal{R}_d$ and $p_h \in \overline{p_\phi}$.

Given a $WDG = \langle N, \rightsquigarrow \rangle$, the instantiation of an ASPIC+ theory AT requires a knowledge base $\mathcal{K} = \mathcal{K}_p$ derived as follows:

- Given $[p_1, \dots, p_q \rightsquigarrow \ell_{Pro} \rightsquigarrow p_\phi], \forall p_i \in \{p_1, \dots, p_q\}$, if p_i is not a conclusion of any inference rule $\bar{A}[\ell_{Pro} \rightsquigarrow p_i] \in WDG, p_i \in \mathcal{K}_p$;
- In addition, assume $WDG' = \langle N', \rightsquigarrow' \rangle$ a subset of WDG containing only a single cycle of inference schemes, then $\forall p_i \in P' \subset N',$ if $[\ell_{Pro_1} \rightsquigarrow p_i], [p_i \rightsquigarrow \ell_{Pro_2}] \in WDG',$ then $p_i \in \mathcal{K}_p$.

The last point is assumed as a method to break cycles among inferences; for example if $[p_1, \dots, p_n \rightsquigarrow \ell_{Pro} \rightsquigarrow p_2], [p_2, \dots, p_m \rightsquigarrow \ell_{Pro} \rightsquigarrow p_3], [p_3, \dots, p_k \rightsquigarrow \ell_{Pro} \rightsquigarrow p_1]$ were included in $WDG', p_1, p_2, p_3 \in \mathcal{K}_p$ in addition to any other premise.

CISpaces.org approach uses a similar mapping to the one adopted between OVA+ [10] and the available solvers (e.g. [20,1]) with the exception of conflict schemes and inference cycles. The mapping semantics of our second type of conflict schemes differs, as we express that $p_1 \wedge \dots \wedge p_n$ is a contrary of p_ϕ while in existing work only one term can be the contrary of another. Premises of inferences that form a cycle in those existing tools are not considered part of the knowledge base.

Given the corresponding ASPIC theory of WDG , we can then compute the preferred extensions and thus highlight the result of the credulous and of the skeptical reasoning process. This differs from systems such as OVA+ [10], as in CISpaces.org different extensions are shown directly in the graph drawn in the interface rather than through external visualisation tools.

Here our key design approach relies on the association of each extension to a hypothesis in the intelligence domain under analysis. Analysts are already familiar with the concept of a hypothesis as explanation of a situation, and are used to elaborate and examine coexistent hypotheses to identify the most plausible one. However, this process normally relies on a matrix approach that may be prone to errors, particularly when hypotheses are interdependent. Thanks to the mapping between the analysis graph and the computation of extensions, analysts are able to easily identify coherent alternative explanations directly on the basis of their analysis.

Figure 3 depicts a possible analysis our analyst would make to assess whether the evacuation road from the Stock Market to the airport via the tunnel is viable. Some of the inference links are also labelled on the basis of the argumentation scheme that they instantiate. Although this is not a feature available in the default installation of

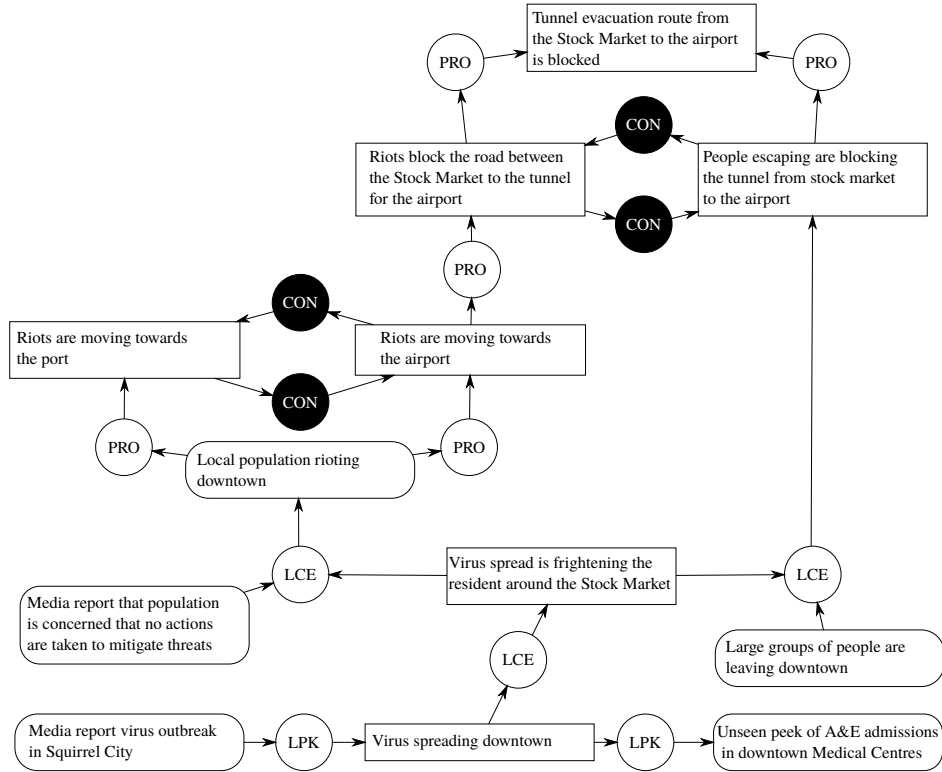


Figure 3. Analysis towards assessing whether the evacuation road from the Stock Market to the airport via the tunnel is a viable one. Rounded boxes represents pieces of information received from open sources; squared boxes claims made by the analyst; white circles inferences (Pro), eventually labelled with the argumentation schemes they refer to; while black circles conflict relations (Con).

CISpaces.org, it is included in the experimental part of the source code: for instance LCE identifies an instance of the scheme for argument from cause to effect (Section 3.1).

It is worth anticipating that this analysis leads to three preferred extensions that can be highlighted to the analysts in terms of the set of acceptable conclusions of arguments. The analyst can clearly add additional pieces of information to change it.

4.3. Desideratum 3: Report Generation

The use of graphical models to represent arguments is the most common approach used in the formal argumentation community to capture argument structures. This requires a significant level of training that cannot be assumed for the recipients of intelligence analysis, viz. decision maker such as group commanders. To this reason, CISpaces.org has been equipped with a Natural Language Generation system. A Natural Language Generation (NLG) system requires [19]: a knowledge source to be used; a communicative goal to be achieved; a user model; and a discourse history. In general, the knowledge source is the information about the domain, while the communicative goal describes the purpose of the text to be generated. The user model is a characterisation of the intended audience, and the discourse history is a model of what has been said so far. A general

NLG system divides processing into a pipeline [19] composed of three stages. First it determines the content and structure of a document (*document planning*); then it looks at syntactic structures and choices of words that need to be used to communicate the content chosen in the document planning (*microplanning*). Finally, it maps the output of the microplanner into text (*realisation*).

In CISpaces.org we followed a rather pragmatic approach regarding the various stages of the NLG pipeline. Indeed, as our main audience are intelligence analysts and decision makers, we strictly obey to the principle of providing them with the important pieces of information in the most concise way. We implemented: (1) a template-based NLG system; (2) a greedy, heuristics-based approach for chaining together premises and conclusion of arguments; (3) an assert-justify writing style suitable for speed reading.

The use of a template-based system is motivated by the fact that our specific audience is used to work with very structured text — cf. for instance the SPOT reports. Our template outputs at the beginning of the produced document the list of statements that are beyond reasonable doubt, i.e. they are conclusions of arguments skeptically accepted according to the preferred semantics. As this report can serve not only as a communication tool when an analyst desires to share their analysis, but also as a summary check for the analyst themselves, we also list all the results of credulous reasoning non-overlapping with the conclusions of skeptically accepted arguments as they are listed at the beginning of the document. Each extension is marked as a plausible hypothesis, cf. Section 4.2.

For each of the skeptically and credulously accepted conclusions we recursively consider the chain of reasoning supporting them, concatenating the various pieces of text with the conjunction *because*. In order not to overwhelm the audience, we choose a greedy heuristic that shows only the final inferential step of each argument at the time, and then recursively expands on its sub-arguments.

In order to support speed reading, which is heavily appreciated by our audience, we consistently employ an assert-justify writing style. Indeed, speed-readers will read the main message as it will always be at the beginning of the line. Moreover, this style informs the reader immediately of the conclusions of arguments. This style is consistent across the entire document, that indeed concludes with the list the pieces of information received, i.e. not generated by the analyst. The rationale for this is that they should be of little value as probably are well-established facts that the analyst uses under the assumption that they have been certified using other techniques [18].

This is the automatically generated report for our running example:

We have reasons to believe that:

- Virus spread is frightening the resident around the Stock Market because Virus spreading downtown
- Virus spreading downtown because [info received] Media report virus outbreak in Squirrel City

Moreover, we also have the following 3 hypotheses.

- Hypothesis number 1
 - * Tunnel evacuation route from the Stock Market to the airport is blocked because Riots block the road between the Stock Market to the tunnel for the airport
 - * Riots block the road between the Stock Market to the tunnel for the airport because Riots are moving towards the airport
 - * Riots are moving towards the airport because [info received] Local population rioting downtown

[...]

Here the pieces of information we received
[...]

It is worth noticing that *Tunnel evacuation route from the Stock Market to the airport is blocked* is a conclusion that is present in each extension, but it is not skeptically accepted as no skeptically accepted argument has it as conclusion.

5. Evaluation and Pathway to Impact

CISpaces.org is the result of a collaboration with the US Army research Lab in the NIS ITA programme and with the UK Defence Science and Technology Laboratory in both the NIS ITA programme and follow-on Defence and Security Accelerator (DASA) programme. CISpaces.org is available for use by professional analysts in both the US (Army Research Laboratory) and the UK (Joint Forces Intelligence Group). The first version of CISpaces [23] was one of three key research highlights in the NIS ITA programme [21]. The refinement of the CISpaces software to take it to Technology Readiness Level 4 (characterised as “validation in a laboratory environment”) was informed by evaluation conducted with professional analysts in the US and the UK as part of the NIS ITA programme, and enabled by the DASA programme. Development work funded by the DASA programme led to CISpaces being made available as an open-source project under a permissive (MIT) licence: <https://github.com/CISpaces>.

Preliminary evaluation of the current interface by a former Royal Signal operator reads: “the system is very good as a demonstrator as a concept.” The operator continues with recommendation to *soldierproof* the interface for “an 18 year old with minimal-to-no schooling may at some point be in charge of this system.”

We also have continued collaboration with the UKs NCA National Cyber Crime Unit, where there is considerable interest in the technologies underpinning CISpaces.org and Open Source Intelligence extraction.

6. Conclusion

CISpaces.org complements human expertise in the generation of reliable intelligence products supporting fact extraction from open sources, mitigating confirmation biases, identifying all possible interpretations of a situation, and explaining why a certain hypothesis is/is not supported by available evidence with the support of an automatically generated report.

The system and its underpinning technologies have attracted positive interest by the UK Joint Forces Intelligence Group as well as by the UK National Crime Agencies Dark Web Intelligence Unit and National Cybercrime Unit.

In future work we will (1) continue the development of argumentation schemes within CISpaces.org; (2) integrate provenance data in the analysis; which will provide us grounds for (3) including preference-based reasoning based also on trust and reputation of sources and information.

References

- [1] Cerutti, F., Vallati, M., Giacomini, M.: An Efficient Java-Based Solver for Abstract Argumentation Frameworks: jArgSemSAT. *International Journal on Artificial Intelligence Tools* 26(02) (2017)
- [2] Chesnevar, C.I., McGinnis, J., Modgil, S., Rahwan, I., Reed, C., Simari, G.R., South, M., Vreeswijk, G.A.W., Willmot, S.: Towards an argument interchange format. *The Knowledge Engineering Review* 21(04), 293 (2006)
- [3] Christensen, J., Mausam, Soderland, S., Etzioni, O.: An analysis of open information extraction based on semantic role labeling. In: *K-CAP 2011*. pp. 113–120 (2011)
- [4] Dung, P.M.: On the Acceptability of Arguments and Its Fundamental Role in Nonmonotonic Reasoning, Logic Programming, and n-Person Games. *Artificial Intelligence* 77(2), 321–357 (1995)
- [5] Dusmanu, M., Cabrio, E., Villata, S.: Argument mining on twitter: arguments, facts and sources. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. pp. 2317–2322 (2017)
- [6] Fader, A., Soderland, S., Etzioni, O.: Identifying relations for open information extraction. In: *EMNLP 2011* (2011)
- [7] Green, N., Ashley, K., Litman, D., Reed, C., Walker, V. (eds.): *Proceedings of the First Workshop on Argumentation Mining*. Association for Computational Linguistics, Baltimore, Maryland (June 2014)
- [8] Heuer, R.: *Psychology of intelligence analysis*. US Government Printing Office (1999)
- [9] IBM: i2 Analyst’s Notebook. <http://www-03.ibm.com/software/products/en/analysts-notebook/>, [Accessed: Feb. 2018]
- [10] Janier, M., Lawrence, J., Reed, C.: OVA+: An argument analysis interface. In: *COMMA 2014*. pp. 463–564 (2014)
- [11] Llewellyn, C., Grover, C., Oberlander, J., Klein, E.: Re-using an argument corpus to aid in the curation of social media collections. In: Calzolari, N., Choukri, K., Declerck, T., Loftsson, H., Maegaard, B., Mariani, J., Moreno, A., Odijk, J., Piperidis, S. (eds.) *LREC*. pp. 462–468 (2014)
- [12] Mausam, Schmitz, M., Bart, R., Soderland, S., Etzioni, O.: Open language learning for information extraction. In: *EMNLP-CoNLL 2012*. pp. 523–534 (2012)
- [13] Middleton, S.E., Middleton, L., Modafferi, S.: Real-time crisis mapping of natural disasters using social media. *IEEE Intelligent Systems* 29(2), 9–17 (2014)
- [14] Middleton, S.E., Krivcovs, V.: Geoparsing and geosemantics for social media: Spatiotemporal grounding of content propagating rumors to support trust and veracity analysis during breaking news. *ACM Transactions on Information Systems* 34(3), 16:1–16:26 (2016)
- [15] Modgil, S., Prakken, H.: A general account of argumentation with preferences. *Artificial Intelligence* 195, 361–397 (2013)
- [16] Prakken, H.: An abstract framework for argumentation with structured arguments. *Argument & Computation* 1(2), 93–124 (2010)
- [17] Procter, R., Vis, F., Voss, A.: Reading the riots on twitter: methodological innovation for the analysis of big data. *International Journal of Social Research Methodology* 16(3), 197–214 (2013)
- [18] Prunckun, H.: *Handbook of Scientific Methods of Inquiry for Intelligence Analysis*. The Scarecrow Press (2010)
- [19] Reiter, E., Dale, R.: *Building Natural Language Generation Systems*. Cambridge (2006)
- [20] Snaith, M., Reed, C.: TOAST: online ASPIC+ implementation. In: *COMMA 2012* (2012)
- [21] The Network and Information Science (NIS) International Technology Alliance (ITA): Online legacy - capstone event. <http://nis-ita.org/capstone> (2016)
- [22] Toniolo, A., Braines, D., Preece, A.D., Webberley, W., Norman, T.J., Sullivan, P., Dropps, T.: Conversational intelligence analysis. In: *Proceedings of the First International Workshop on Understanding Situations Through Multimodal Sensing*. pp. 42:1–42:6 (2016)
- [23] Toniolo, A., Norman, T.J., Etuk, A., Cerutti, F., Ouyang, R.W., Srivastava, M., Oren, N., Dropps, T., Allen, J.A., Sullivan, P.: Supporting Reasoning with Different Types of Evidence in Intelligence Analysis. In: *AAMAS 2015*. pp. 781–789 (2015)
- [24] UK DSTL Defence and Security Accelerator: Scenario: Non-Combatant Evacuation Operation (NEO). <https://tinyurl.com/y9kywro6>, [Accessed: Feb. 2018]
- [25] Walton, D., Reed, C., Macagno, F.: *Argumentation schemes*. Cambridge University Press, NY (2008)
- [26] Wu, F., Weld, D.S.: Open information extraction using wikipedia. In: *ACL 2010*. pp. 118–127 (2010)