Making Informed Decisions with Provenance and Argumentation Schemes

Alice Toniolo[1], Federico Cerutti[1], Nir Oren[1], Timothy J. Norman[1], and Katia Sycara[2]

[1] Department of Computing Science, University of Aberdeen, Scotland, UK
{a.toniolo,f.cerutti,n.oren,t.j.norman}@abdn.ac.uk
[2] Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, US
katia@cs.cmu.edu

Abstract. Intelligence analysis is the process of reasoning about information in order to produce hypothetical explanations for a situation. In this process, it is fundamental to assess how, when, and where this information has been elaborated. A model of provenance can capture this contextual information. Provenance data inevitably affects the identification of plausible conclusions, thus, it must be introduced in the reasoning process. In this paper, we propose a model of argument schemes that allows software agents to explore provenance for improving the information assessment. Argument schemes present the essential elements of provenance that warrant the credibility of the information. Schemes are also used to establish preferences between pieces of information according to different provenance criteria, such as timeliness and reliability. The introduction of schemes about provenance facilitates the decision-making process by providing a rational method to assess the credibility of a piece of information and to resolve conflicting information.

Keywords: Argumentation schemes, information quality, provenance

1 Introduction

When receiving conflicting information on a situation of interest, an agent has to make a decision on which piece of information should be considered. In order to reach a decision, an agent may look for further information to clarify the event. Alternatively, the agent may consider additional information about how, when, and where the information was gathered in the formation of the claim. Provenance is a rigorous way to capture this contextual information [4, 8]. Reasoning about information provenance lies at the heart of intelligence analysis, where the aim of analysts is to gather evidence for constructing plausible hypotheses for explaining a situation [11]. An analyst must consult provenance data in assessing the reliability of information according to timeliness, trustworthiness and so on. More importantly, an analyst must be aware of this assessment when identifying plausible conclusions. This may lead to further exploration of the provenance data to understand how a piece of information was handled and why. In the light of new provenance information, the belief about the plausibility of certain conclusions may change. This paper, therefore, aims to address the question of how we can provide support to analysts in evaluating hypotheses while exploring and assessing the provenance of their supporting evidence.

Argumentation techniques, thanks to their ability of encompassing non-monotonic reasoning, appear to be useful to addressing this issue. In particular, argumentation schemes have been applied to evidential reasoning for representing patterns to model the warrants underpinning evidential inferences [1]. Critical questions, then, focus on the exploration of further evidence for a claim. To date, however, formalised models of provenance within argumentation-based reasoning are yet to be explored. The *argument scheme from position to know* is a way to introduce the provenance of an assertion into reasoning [15]. However, it is assumed that the assertion has been directly stated by the source. In intelligence analysis, information is manipulated many times before being analysed, thus, the entire elaboration process must be recorded as a provenance chain, and introduced into reasoning.

In this paper, we propose a model of argumentation schemes that allows agents to reason and explore provenance chains while drawing and assessing conclusions. We discuss how such schemes may be formalised in a defeasible language in order to be handled in an argumentation framework. Provenance is a novel field of application for argumentation-based systems. A similar direction has been investigated in [2]. This work records provenance of arguments only to be consulted for post-analysis. In contrast, here we include provenance as part of the reasoning process. Moreover, we propose a scheme used to verify the quality of statements and select those more preferred because they are retrieved with a more reliable, accurate or trustworthy process. In existing research, preferences have been widely discussed within argumentation frameworks [7, 13], however, how to construct them is still an open question [3]. We address this problem by defining criteria based upon the analysis of provenance information. We believe that by exploring provenance data, an agent will be able to support analysts in making more informed decisions for the identification of plausible hypotheses.

2 Argumentation Schemes and Provenance

In this section, we discuss our contribution towards the development of argumentation schemes for provenance elements and for preferences derived from provenance criteria. Retrieving information in the context of intelligence analysis is a complex process that may involve many steps of elaboration before the analyst is able to analyse it. Furthermore, analysts have to deal with a large amount of collected data, which may contain noise and misleading information. In this research, we argue that rigorously recording contextual data, through the use of a formalised provenance model, and reasoning about it will facilitate the assessment of the credibility of the information, reducing the analytical effort in sense-making.

The underpinning language for provenance is the W3C standard PROV Data

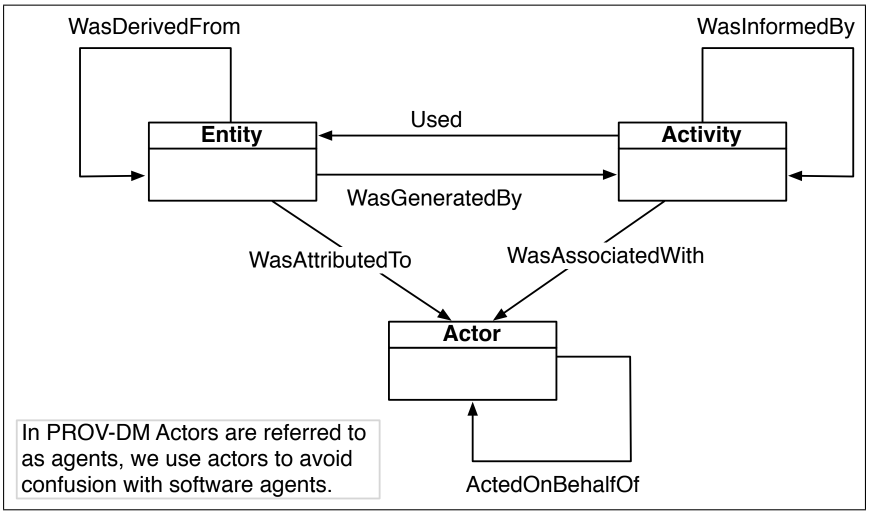

Fig. 1. The PROV-DM core

Model (PROV-DM[8]). PROV-DM records provenance in terms of *entities*, *activities*, and *actors* that have caused an entity to be and it defines seven relationships between these elements (Figure 1). An entity is a physical or conceptual thing such as a report or a piece of information; an entity may be derived from other entities. An activity represents a process that acts upon entities; e.g., extracting, creating entities. Entities are generated by an activity, and they represent resources that can be consumed (used) by other activities. An activity may inform another activity by triggering it to take place. An actor is something or someone responsible for an activity taking place such as a person, or a software tool. An actor may author an entity or it may act on behalf of other actors.

In order to include provenance in the process of intelligence analysis, we rely on the concept of argumentation schemes [15]. 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) that can be used to challenge the validity of arguments. Using these schemes, conclusions can be assessed in terms of the evidence gathered, in particular challenging assumptions about its credibility warranted by the provenance of the information.

We illustrate our model by considering the example in Figure 2. The goal for analyst Joe is to establish whether criminal activity is present on the northern border of a region. A gang G is suspected of smuggling forbidden products across this zone. Joe wants to understand if gang G has crossed the border. An informant discloses that gang G has crossed the North border (statement i_k). This information was sent to the informant from an observer, via a messenger. Joe also receives a piece of information from the surveillance cameras stating that gang G is moving south (statement i_j). Joe may, then, infer that the gang

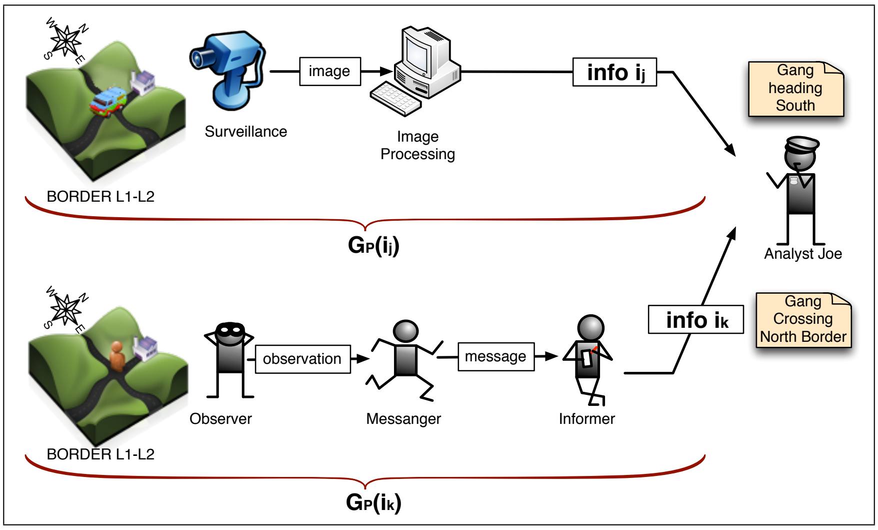

Fig. 2. Scenario example

has not crossed the North border. Since there is a conflict of information, Joe has to decide whether gang G has, or has not, crossed the border. An option would be to verify the primary sources; i.e., consulting images or the original observer. However, since analysts may not have access to these sources, the decision must be made only according to the provenance information available.

In this context, Joe has to assess whether an item of information is credible and decide which one is preferred in order to make a decision. Joe analyses the credibility of i_k by considering its source, the observer O. Using an *argumentation scheme from position to know* [15], Joe would state that:

- Observer O is in a position to know whether i_k is true,
- Observer O asserts that i_k is true,

$\Rightarrow$ Therefore, i_k may plausibly be taken to be true.

The critical questions are the following. *Group 1:* Is O in a position to know whether i_k is true? Is O an honest/trustworthy source? *Group 2:* Did O assert that i_k is true? Is i_k consistent with reports from other sources? Is i_k supported by evidence? The first group of critical questions challenges assumptions on the position, the credibility and the role of the observer O. The second group challenges the assertion. However, the argument from position to know does not permit an agent to state, for example, that the messenger, not the observer, was untrustworthy, or that the informer has reported an incorrect observation. In the next section, we address the need of providing adequate schemes for provenance.

2.1 An Argumentation Scheme for Provenance

Here we introduce an argumentation scheme that permits an agent to consider partial provenance data in evaluating a piece of information. We informally refer to the provenance chain as the *chain of elaboration* $G_P(i_j), G_P(i_k)$ that includes activities, actors and entities, and that lead from the primary source (the observer and the camera in Figure 2) to the information in hands of the analyst, i_j and i_k. This will be formally defined in Section 3.

This scheme extends the argumentation scheme from position to know. There are two main statements at the base of this scheme, that lead an agent to believe that a statement i_j holds. The first is about observer O stating that i_j is true. The second is a warrant for the assertion i_j. The fact that O is in a position to know about i_j, functions as a warrant for i_j being plausibly true. We adopt a similar structure by assuming that the provenance chain is a warrant for stating that a piece of information i_j is credible. The critical questions will be used to gradually explore relevant provenance data, in particular when a conflict between statements arises. We extract the important elements of the provenance chain using the model of the *7Ws of provenance* [4]. This model is based upon the 7 questions "Who?, What?, Where?, Why?, Which?, When?", and "How?" about events that affect data during its lifetime. Important elements of the provenance chain (such as derivation from documents, goal of the analysis and so on) form patterns of relationships as those in PROV-DM (Figure 1) that are extracted using the 7Ws model and introduced into the scheme. Here, we informally define the argumentation scheme for provenance Arg_A, we will formalise these concepts in the forthcoming sections.

Definition 1 (Argumentation scheme for provenance Arg_A).

- *Given a provenance chain $G_P(i_j)$ of i_j, information i_j:*
  - *(Where?) was derived from an entity A*
  - *(Who?) was associated with actor AG*
  - *(What?) was generated by activity P1*
  - *(How?) was informed by activity P2*
  - *(Why?) was generated to satisfy goal X*
  - *(When?) was generated at time T*
  - *(Which?) was generated by using some entities A1,..., AN*
  - *where A, AG, P1, ... belong to $G_P(i_j)$*
- *the stated elements of $G_P(i_j)$ ensure that information i_j is true,*

$\Rightarrow$ *Therefore, information i_j may plausibly be taken to be true.*

The core statements of this scheme are: the derivation of information node i_j from the provenance chain $G_P(i_j)$ and the fact that the provenance chain warrants the credibility of information i_j. The critical questions that we can define for the former statement focus on further reasons, unrelated to the provenance of i_j, for believing that i_j is plausible. These questions link this scheme with the analysis of intelligence in the formation of hypotheses. The questions are:

- **CQA1:** Is i_j consistent with other information?
- **CQA2:** Is i_j supported by evidence?

Question CQA1 drives a search for further information. Some conditions can be added to this question to capture a stronger conflict. For example, we may ask whether i_j is consistent with an information item that shares part of the provenance chain, such as the derivation from the same sources A1,..., AN.

Other CQs may challenge the credibility of i_j warranted by provenance:

- **CQA3:** Does $G_P(i_j)$ contain other elements that lead us not to believe i_j?
- **CQA4:** Is there any other provenance element that should have been included to ensure that i_j is believable?

These critical questions ask for more provenance data to verify that, even in the light of this new data, i_j is still credible. Using CQA3 an agent searches for patterns that specifically deny the credibility of i_j. Question CQA4 searches for missing patterns that should have been included to warrant the credibility of i_j.

An agent may verify the credibility of i_j with argument Arg_A. Nevertheless, information may be conflicting and agents must make a decision on which statement is preferred. In our example, we discussed the need for analyst Joe to decide if the gang G has crossed the border, when two reports are conflicting. We now describe how this decision can be made using provenance preferences.

2.2 An Argumentation Scheme for Provenance Preferences

Here we provide an argument scheme that expresses a preference between two pieces of information i_j and i_k in order to enable the resolution of conflicts. The scheme states a reason for a preference to exist, based upon some dimensions of quality (similar to the preferences in [3]). Several dimensions have been proposed to test the quality of information such as timeliness, relevancy, trustworthiness [10, 5]. There are different ways to perform assessments according to these dimensions, as information quality is often dependent upon the domain [10]. For the purpose of this research, we specify criteria to identify preferences between information items, which can be tested using provenance elements and are relevant in the context of intelligence analysis [14]. These are:

Trustworthiness: a piece of information is trustworthy if all the sources of that information report truthfully. A provenance chain of an information item may be more trustworthy than another if its least trustworthy actor is more trustworthy than that of the other chain. The motivation is that if there was at least an untrustworthy actor elaborating the information, a very skeptical analyst would also consider that information untrustworthy.

Reliability: a piece of information is reliable if every activity that has contributed to the generation of that information performs consistently well. A provenance chain may be more preferred than another if its least reliable process is more reliable than that of the other chain.

Timeliness: a piece of information is timely if it is more recent. Given two information entities, if the latter is older than the former, an agent may consider the former preferred to the latter.

Accuracy of derivation: a piece of information is more accurate if it comes directly from its primary source. The shortest path of derivation from the furthest primary source to the statement may indicate a preference over competing information. In fact, an information item may be considered less accurate if it has been elaborated many times before being processed.

Using these criteria, an agent may state that the piece of information belonging to the preferred provenance record is preferred to other statements. Here we informally define the essential elements of our framework for expressing preference criteria based on provenance. We refer to a general preference between information with symbol $\ll$. If the preference is stated according to some criteria or function β, it is represented with $\ll_\beta$. Note that $a \ll b$ may or may not correspond to $a \ll_\beta b$. A priority between inference rules is represented with $\prec$, while $<$ is the usual maths comparison operator.

Let us indicate the stated part of a provenance chain $G_P(i_j)$ with $\overline{G}_P(i_j)$. A preference is based on a criterion Crt, and it is defined between elements of two provenance chains $\overline{G}_P(i_j) \ll_{Crt} \overline{G}_P(i_k)$, where $\overline{G}_P(i_k)$ is preferred to $\overline{G}_P(i_j)$. The argument Arg_B states that if this is the case, then an agent can infer that information i_k is preferred to i_j, $i_j \ll i_k$. The argument is:

Definition 2 (Informal argument for provenance preferences Arg_B).
- *Given information i_j and i_k,*
- *and their known parts of the provenance chains $\overline{G}_P(i_j)$ and $\overline{G}_P(i_k)$,*
- *if there exists a criterion Ctr such that $\overline{G}_P(i_j) \ll_{Ctr} \overline{G}_P(i_k)$, then $i_j \ll i_k$.*
- *a criterion Ctr' leads to assert that $\overline{G}_P(i_j) \ll_{Ctr'} \overline{G}_P(i_k)$*

$\Rightarrow$ *Therefore, i_k should be preferred to i_j.*

The critical questions for this argument are:

- **CQB1:** Does a different criterion Ctr_1, such that $\overline{G}_P(i_j) \gg_{Ctr_1} \overline{G}_P(i_k)$ lead $i_j \ll i_k$ not being valid?
- **CQB2:** Is there any exception to criterion Ctr such that even if a provenance chain $\overline{G}_P(i_k)$ is preferred to $\overline{G}_P(i_j)$, information i_k is not preferred to information i_j?
- **CQB3:** Is there any other reason for believing that the preference $i_j \ll i_k$ is not valid?

This scheme is used by agents to decide whether an information item has a greater quality than others and thus it should be preferred.

We presented two argumentation schemes for exploring provenance in the reasoning process. The first claims that an agent's statement is justified by its provenance record, and so it is credible. Further, an agent may prefer one piece of information to another if some provenance elements of the former are preferred to some elements of the latter. We now define a means to capture provenance records for constructing the sorts of arguments we are interested in.

3 Provenance

How may we formally record and query provenance elements for instantiating our argumentation schemes? As previously introduced, our underpinning language is PROV-DM [8], which defines provenance in terms of agents, entities, and activities involved in producing, influencing, or delivering an entity. A record of provenance is formed by nodes (entities, actors, activities) and directed relationships between these nodes. Such a record can be represented as a directed acyclic graph. We may then explore these graphs using OPQL [6], a provenance query language that supports lineage queries. Here we present our extension of OPQL for dealing with PROV-DM.

Definition 3 (Provenance Graph). *A provenance graph is graph* $G_P = (N, E)$ *that consists of:*

- *a set of nodes* $N = A \cup AC \cup AG = \{n_1, n_2, \dots\}$ *and:*
  - A *is a set of* entities, $A = \{a_1, a_2, \dots\}$
  - AC *is a set of* activities, $AC = \{p_1, p_2, \dots\}$
  - AG *is a set of* actors, $AG = \{ag_1, ag_2, \dots\}$
- *a set of directed edges* $E = E_u \cup E_g \cup E_d \cup E_i \cup E_{aw} \cup E_{at} \cup E_b$ *where*
  - $E_u \subseteq AC \times A$ *with* $(p, a) \in E_u$*: activity* p used a,
  - $E_g \subseteq A \times AC$ *with* $(a, p) \in E_g$*: entity* a was generated by p,
  - $E_d \subseteq A \times A$ *with* $(a_1, a_2) \in E_d$*: entity* a_1 was derived by a_2,
  - $E_i \subseteq AC \times AC$ *with* $(p_1, p_2) \in E_i$*: activity* p_1 was informed by p_2,
  - $E_{aw} \subseteq AC \times AG$ *with* $(p, ag) \in E_{aw}$*: activity* p was associated with ag,
  - $E_{at} \subseteq A \times AG$ *with* $(a, ag) \in E_{at}$*: entity* a was attributed to ag,
  - $E_b \subseteq AG \times AG$ *with* $(ag_1, ag_2) \in E_b$*: actor* ag_1 acted on behalf of ag_2.

Note that a node n refers to a general node of type entity a, activity p, or actor ag. Nodes n and edges e comprise a set of attribute-value pairs. Given a set of attributes $Att = \{attribute_1, attribute_2, \dots\}$ and a set of corresponding values $Val = \{value_1, value_2, \dots\}$, a mapping function $att : E \cup N \times Att \rightarrow Val$ associates a value to an attribute of an edge or a node. For example, the name Inf_1 of an entity a_1 is $att(a_1, name) =$ "Inf_1", the time associated with a generation edge $e_1 = (a, p)$ is $att(e_1, time) =$"2014-01-22:T11-51-00".

In our system, an agent has two datasets available $\mathcal{I}$ and $\mathcal{P}$. The dataset $\mathcal{I} = \{\dots\}$ includes pieces of information referred to as nodes i_j. $\mathcal{P}$ contains a graph of provenance data for information nodes in $\mathcal{I}$. Here we recall some graph properties to represent a provenance chain.

Definition 4 (Graph union). *The* union *of two subgraphs is represented as* $G_{P1} \cup G_{P2}$ *whereby* $G_{P1} = (N_1, E_1)$ *and* $G_{P2} = (N_2, E_2)$ *and it results in a new provenance graph* $G'_P = (N', E')$ *where:*

$$N' = \{n | n \in N_1 \cup n \in N_2\} \qquad E' = \{e | e \in E_1 \cup e \in E_2\}$$

Definition 5. *A* directed path *from a node* n_0 *to a node* n_k *is a directed graph* $D_P(n_0, n_k) = (N, E)$ *with distinct nodes* $N = \{n_0, \dots, n_k\}$ *and edges* $E = \{e_0, \dots, e_{k-1}\}$ *such that* e_i *is an edge directed from* n_i *to* n_{i+1}*, for all* $i < k$.

There exists a number of directed paths $D_{Pi}(n_0, n_k) = (N_i, E_i)$ *between the two nodes* n_0 *and* n_k*. The length of a path is the cardinality of the edge set* $|E_i|$.

The shortest directed path $S_P(n_0, n_k)$ *is the path where the cardinality of the edge set is the minimum.*

$$\begin{aligned} &Given \quad D_{P1}(n_0,n_k) = (N_1,E_1),\dots,D_{Pu}(n_0,n_k) = (N_u,E_u) \\ &S_P(n_0,n_k) = D_{Pq}(n_0,n_k) = (N_q,E_q) \quad where \quad E_q = min_{i=1,\dots,u}(|E_i|) \end{aligned}$$

Definition 6 (Provenance chain). *A provenance chain of a node n_j in $\mathcal{P}$ is a subgraph $G_P(n_j) = (N', E')$ of $G_P = (N, E)$*

$$G_P(n_j) = \bigcup_{n_q \in N : \exists D_P(n_j,n_q), \nexists n_l(n_q,n_l) \in E} D_P(n_j, n_q)$$

Graph $G_P(n_j)$ represents a union between all the paths from node n_j in $\mathcal{P}$ to a node $n_q \in N$ that does not have successors. Note that in Section 2 we informally introduced $G_P(i_j)$ as the provenance chain of information i_j. In more formal terms, the provenance chain $G_P(i_j)$ indicates a graph $G_P(n_j)$ of an entity node n_j that is linked to information i_j through $att(n_j, name) = i_j$ (eg., i_j= "The gang is heading south", $att(n_j, name)$ = "The gang is heading south"). Henceforth, for convenience we will refer to $G_P(i_j)$ in general discussion, but the formalisation is presented in terms of the correspondent graph $G_P(n_j)$.

Given a provenance graph $G_P(n_j)$, a query to the provenance dataset $\mathcal{P}$ in OPQL is made by using graph patterns and pattern matching.

Definition 7. *A* graph pattern *is a pair $P_m = (G_M, C)$, where $G_M = (N_M, E_M)$ is a graph motif and C is a predicate*[3] *on the attributes of the motif. A* graph motif *G_M is a graph with a certain structure but where nodes and edges are identified by a variable.*

A graph pattern $P_m = (G_M, C)$ is matched *with a graph $G_P = (N, E)$ if there exists an injective mapping $\phi : N_M \rightarrow N$ such that:*

i) $\forall e(n_1, n_2) \in E_M$, the mapping $(\phi(n_1), \phi(n_2))$ is an edge in $E \in G_P$
ii) predicate C holds in the mapping of G_M in G_P

The matched graph is a graph identified by $\langle \phi, P_m, G_P \rangle$ and referred to as $\phi_{P_m}[G_P]$.

A graph pattern is a variable that permits the extraction of the structure required by the pattern. An example of a 1-node pattern that extracts all nodes that are labelled "Observer" is:

$$\begin{aligned} &P_b = (G_M, C) \qquad G_M = (N_M, E_M) \\ &N_M = \{n_1\} \qquad E_M = \emptyset \qquad C : att(n_1, name) = \text{"}Observer\text{"} \end{aligned}$$

A 2-node pattern extracts an edge between two nodes. These 1-node or 2-node patterns are used to perform queries in order to extract a named node or a named edge with specific attributes.

In Figure 3, a formal version of our example (Figure 2) showcases the use of PROV-DM for provenance chains. In the figures, we use the PROV-DM notation to refer to an agent as a pentagon, to an entity as a circle and to an activity as a squared box. A provenance graph reads right to left, whereby activities towards the lefthand side are older than those towards the righthand side of the graph.

[3] Intuitively, C is similar to the SQL condition "WHERE" in a "SELECT" query.

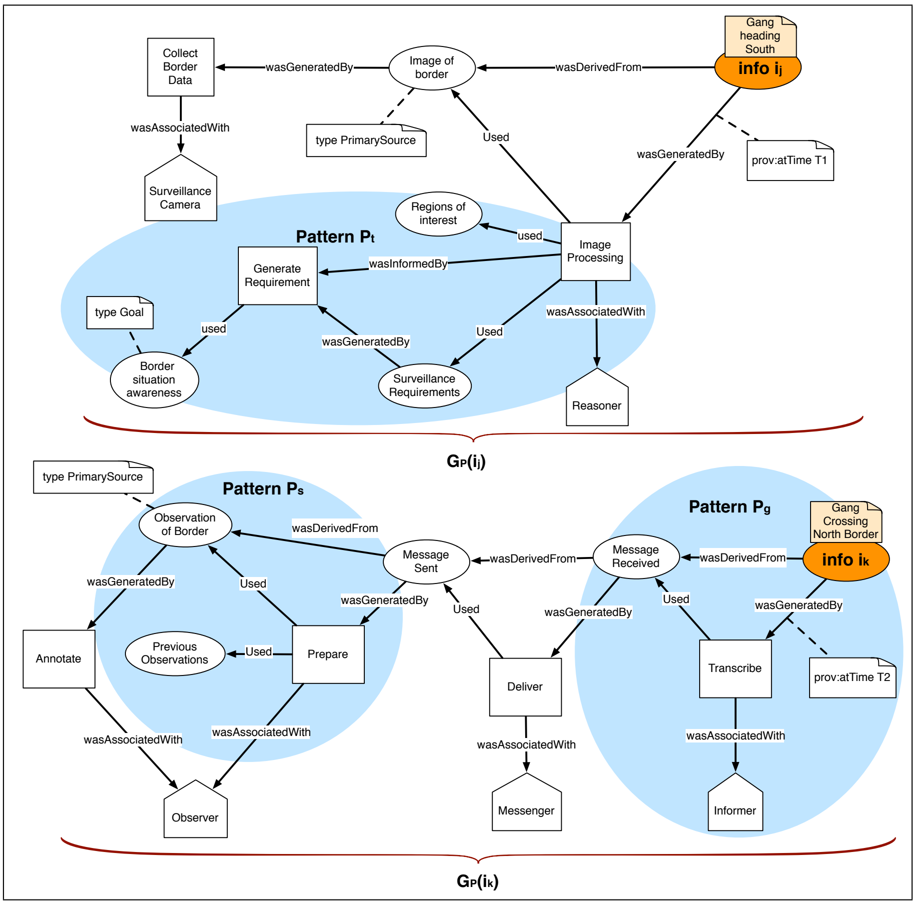

Fig. 3. PROV-DM model of the provenance chain in Fig.2

4 Provenance Elements for Argument Arg_A

Now we may specify a procedure to query provenance data for instantiating an argument Arg_A. In intelligence analysis, the dataset of information $\mathcal{I}$ is extracted from documents, reports, newspapers, emails, and so on. We observed that the provenance of this type of data presents recurrent patterns P_m. These can be then used to extract data from the provenance dataset $\mathcal{P}$. Here we formally define these patterns, and we express them using the 7W-questions [4] within the scheme for provenance Arg_A. The steps for instantiating Arg_A are:

1. pattern variables P_m are used to query the provenance chain $G_P(n_j)$ in $\mathcal{P}$ and extract frequent patterns;
2. the result of the query, the matched graph $\phi_{P_m}[G_P(n_j)]$, is used to generate a 7W-explanation of provenance;

3. the 7W-explanation of $\phi_{P_m}[G_P(n_j)]$ is introduced in the argument, to support the conclusion that i_j is plausibly true.

A *7W-explanation* is a set of at most 7 statements that describe a pattern P_m in the provenance graph $G_P(i_j)$. Each statement w_i represents a linguistic explanation of a provenance relationship between the node n_j linked to information i_j and another node, edge or attribute that is related to n_j through the matched graph $\phi_{P_m}[G_P(n_j)]$. This is extracted by answering one of the 7W-questions of provenance. In this research, we define the following correspondence to derive the 7W-explanations of the provenance chain $G_P(n_j)$:

- "What?" denotes an activity p_1 that generated an entity a_1
- "When?" refers to the time at which the generation of a_1 occurred
- "Where?" generally refers to a physical location, but we consider here the source of where the entity a_1 has a direct derivation edge with a_2
- "How?" is the activity p_2 that informed the activity p_1
- "Who?" refers to the actors associated with the activity p_1
- "Which?" denotes the sources or instruments used by activity p_1 that are not linked to the entity a_1 with a "wasDerivedFrom" edge.
- "Why?" is the reason for p_1, we report an entity a_2, marked with attribute $att(a_2, type) =$ "*Goal*"

For example, in Figure 3 assume that, using a 2-node pattern, we extract the derivation edge of information i_k stating that i_k was derived from the received message. The 7W-explanation of this pattern contains only one statement $w_1 =$"i_k was derived from Message Received" corresponding to the question "Where?".

4.1 Provenance Patterns for Intelligence Analysis

Here, we describe frequent patterns P_m for intelligence analysis. Their correspondence to the 7W-questions is highlighted in the figures.

Extraction of information and updates. We define a pattern P_g for generating entities, with the following characteristics:

$$P_g = (G_M, C) \qquad G_M = (N_M, E_M)$$
$$N_M = \{a_1, a_2, p_1, ag_1\} \qquad E_M = \{e_1, e_2, e_3, e_4\}$$
with $e_1 = (a_1, a_2) \in E_d, e_2 = (a_1, p_1) \in E_g,$
$e_3 = (p_1, a_2) \in E_u, e_4 = (p_1, ag_1) \in E_{aw}$

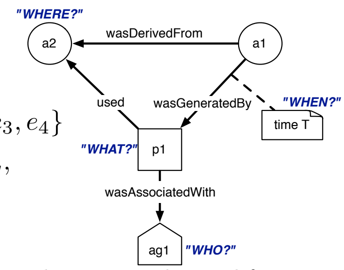

This pattern takes two entities, a_1 and a_2, whereby a_1 was derived from a_2. Activity p_1 was responsible for generating entity a_1 using a_2 and it was associated with actor ag_1. This pattern is the most frequent and it can be used, for example, to identify an extraction of an information node with $C : att(p_1, name) =$"Extract".

Preparation of a document and primary sources. The preparation of a document is a source pattern P_s:

$$P_s = (G_M, C) \qquad G_M = (N_M, E_M)$$
$$N_M = \{a_1, a_2, \ldots, a_n, p_1, ag_1\}$$
$$E_M = \{e_1, e_2, e_3, \ldots, e_n\} \text{ with}$$
$$e_1 = (a_1, p_1) \in E_g, e_2 = (p_1, ag_1) \in E_{aw},$$
$$e_3 = (p_1, a_2), \ldots, e_n = (p_1, a_n) \in E_u$$

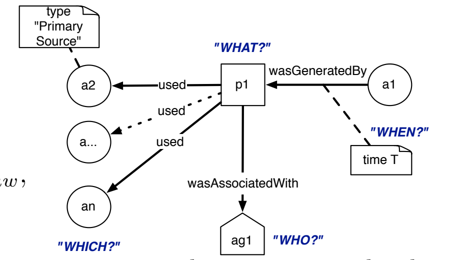

The centre of the provenance record is an activity p_1 that generates the document recorded in entity a_1 and uses a number of sources $a_2, \ldots, a_n$. An important attribute qualifies an entity as the primary source, where $att(a, type) =$"Primary Source". Primary sources are those that first reported or created the information. In this way, we can reason about how far an item of information is from originating sources.

Intelligence requirement or goal of analysis. Here we define a pattern P_t fundamental for recognising the goal of the analysis. This may also be called an intelligence requirement or a request for information. The pattern P_t denotes the triggering activity p_2 that caused activity p_1 to be executed. Goals are marked with attribute $C : att(a_3, type) =$ "*Goal*".

$$P_t = (G_M, C) \quad G_M = (N_M, E_M)$$
$$N_M = \{a_1, a_2, a_3, p_1, p_2\}$$
$$E_M = \{e_1, e_2, e_3, e_4, e_5\}$$
$$e_1 = (a_1, p_1), e_3 = (a_3, p_1) \in E_g,$$

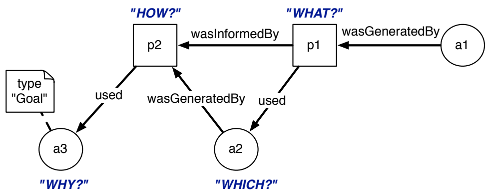

and $e_2 = (p_1, a_2), e_4 = (p_2, a_3) \in E_u, e_5 = (p_1, p_2) \in E_i$.

Using the above patterns, agents can construct arguments for provenance, Arg_A.

5 Preference Criteria for Argument Arg_B

In this section, we define how to compare elements of a provenance chain to reason about whether one piece of information is preferred to another. The criteria for intelligence analysis presented in Section 2.2 are used to instantiate argument scheme Arg_B.

We recall that a provenance chain $G_P(n_k)$ preferred to $G_P(n_j)$ is referred to as $G_P(n_j) \ll_{Ctr} G_P(n_k)$. While in argument Arg_B the preferences stated are established between the known part of the graphs $\overline{G}_P(n_k)$ and $\overline{G}_P(i_j)$, we assume here that an agent knows the whole chains $G_P(n_j)$ and $G_P(n_k)$ in order to simplify the notation. In the argument Arg_B, it is assumed that if the preference $G_P(n_j) \ll_{Ctr} G_P(n_k)$ is valid, then the information i_k belonging to the provenance chain $G_P(n_k)$ is also preferred to the information i_j, $i_j \ll i_k$.

Definition 8 (Criterion Ctr). *A criterion Ctr to compare provenance chains $G_P(n_j)$ and $G_P(n_k)$, stating whether $G_P(n_j)$ is preferred to $G_P(n_k)$, is formed by a tuple $Ctr = \langle Ptype, Cond \rangle$ where Ctr indicates the dimension of the quality assessment, Ptype indicates what type of pattern/path should be compared for that dimension, and Cond states how to compare them.*

The steps used to instantiate Arg_B are:

1. A criterion $Ctr = \langle Ptype, Cond \rangle$ is chosen to compare parts of provenance chains $G_P(n_j)$ and $G_P(n_k)$ for nodes n_j and n_k respectively.
2. *Ptype* is used to query the provenance chains $G_P(n_j)$ and $G_P(n_k)$ and extract the relevant elements to be compared. *Ptype* is:
   T1. a pattern P_m matched against both provenance chains, resulting in matched graphs $\phi_{P_m}[G_P(n_j)]$ and $\phi_{P_m}[G_P(n_k)]$;
   T2. directed paths $D_{P1}(n_j, n_a)$ and $D_{P2}(n_k, n_b)$ to nodes $n_a \in G_P(n_j)$ and $n_b \in G_P(n_k)$ extracted through simple 1-node patterns P_b.
3. The output of the query using *Ptype* always returns two partial graphs that we call $G_Q(n_j) \subseteq G_P(n_j)$ and $G_Q(n_k) \subseteq G_P(n_k)$.
4. *Cond* specifies how to define a preference between the elements extracted from the above query. *Cond* is specific for the criterion used. *Cond* establishes that $q_j \in G_Q(n_j) \ll_{Cond} q_k \in G_Q(n_k)$, where q_i may be:
   K1. an attribute $att(n, attribute)$ of a node n
   K2. an attribute $att(e, attribute)$ of an edge e
   K3. a count on the elements of $G_Q(n_i)$
5. Given a condition *cond* that enforces a preference between $q_j \in G_Q(n_j) \ll_{Cond} q_k \in G_Q(n_k)$, we state that graph $G_P(n_j)$ is less preferred than $G_P(n_k)$, $G_P(n_j) \ll_{Ctr} G_P(n_k)$.
6. The result of the preference $G_P(n_j) \ll_{Ctr} G_P(n_k)$ is introduced in the argument Arg_B to support the preference between i_j and i_k, $i_j \ll i_k$.

This general procedure permits us to state preferences for the intelligence analysis criteria presented in Section 2.2. Here we briefly introduce them and in Table 1 a more detailed procedure is proposed.

Trustworthiness and Reliability. The preference is asserted on the basis of the least trustworthy actor in each provenance chain, $ag_j \in G_P(n_j)$ and $ag_k \in G_P(n_k)$. An overall order of actor names is defined by $Cond$, from which we can extract, for example, $ag_j \ll_{Cond} ag_k$. In the same way, reliability is assessed on the basis of the least reliable activity and an overall order of activity names.

Timeliness. We extract all the generation edges, where the attribute time is defined, for each of the provenance chains. The assessment is based on the most recent timestamp among each set of edges; e.g., T_j of $G_P(n_j)$ and T_k of $G_P(n_k)$. Time is ordered, thus, when $T_j < T_k$ the condition also states $T_j \ll_{Cond} T_k$.

Accuracy of derivation. In this case, we extract all the shortest directed paths from n_j and n_k to their primary source nodes. The longest of these paths for each node n_j and n_k is used for assessment; i.e., the furthest of the primary sources. The shortest path between the two paths is the preferred one.

Trustworthiness ($q_i = K1$, $Ptype = T1$) **Input -** list of actor names $UAG = \{actor_1, actor_2, \dots\}$, **Preferred order -** most trustworthy: $actor_1 \gg \dots \gg actor_n$: least trustworthy **Ptype -** $P_b = (G_M, C)$, $N_M = \{ag\}$, $G_Q(n_j) = \phi_{P_b}[G_P(n_j)] = \{ag_a, \dots, ag_l\}$, $G_Q(n_k) = \phi_{P_b}[G_P(n_k)] = \{ag_b, \dots, ag_m\}$ where for each ag_ℓ there exists $att(ag_\ell, name) = actor_\ell$. **Cond -** Given the least trustworthy actors ag_j and ag_k: - ag_j: $att(ag_j, name) = actor_j$, $actor_j \ll actor_i$, $actor_i = att(ag_i, name)$ for all $i = a, \dots, l$ - ag_k: $att(ag_k, name) = actor_k$, $actor_k \ll actor_i$, $actor_i = att(ag_i, name)$ for all $i = b, \dots, m$ - The most trustworthy is preferred: if $actor_j \ll actor_k$ then $ag_j \ll_{Cond} ag_k$ **Conclusion -** $ag_j \ll_{Cond} ag_k$, then $G_P(n_j) \ll_{Trust} G_P(n_k)$
Reliability ($q_i = K1$, $Ptype = T1$) **Input -** list of activity names $UPR = \{activ_1, activ_2, \dots\}$, **Preferred order -** most reliable: $activ_1 \gg \dots \gg activ_n$: least reliable **Ptype -** $P_b = (G_M, C)$, $N_M = \{p\}$, $G_Q(n_j) = \{p_a, \dots, p_l\}$, $G_Q(n_k) = \{p_b, \dots, p_m\}$ where for each p_ℓ there exists $att(p_\ell, name) = activ_\ell$. **Cond -** Given the least reliable activities p_j and p_k : -p_j: $att(p_j, name) = activ_j$, $activ_j \ll activ_i$, $activ_i = att(p_i, name)$ for all $i = a, \dots, l$ -p_k: $att(p_k, name) = activ_k$, $activ_k \ll activ_i$, $activ_i = att(p_i, name)$ for all $i = b, \dots, m$ - The most reliable is preferred: if $activ_j \ll activ_k$ then $p_j \ll_{Cond} p_k$ **Conclusion -** $p_j \ll_{Cond} p_k$, then $G_P(n_j) \ll_{Rel} G_P(n_k)$
Timeliness ($q_i = K2$, $Ptype = T1$) **Ptype -** $P_e = (G_M, C)$, $N_M = \{a, p\}$, $E_M = \{e = (a, p) \in E_g\}$, $G_Q(n_j) = \{e_a, \dots, e_l\}$, $G_Q(n_k) = \{e_b, \dots, e_m\}$ where for each e_ℓ there exists $att(e_\ell, time) = T_\ell$. **Cond -** Given the most recent times T_j and T_k: - T_j: $att(e_j, time) = T_j$, $T_j > T_i$, $T_i = att(e_i, time)$ for all $i = a, \dots, l$ - T_k: $att(e_k, time) = T_k$, $T_k > T_i$, $T_i = att(e_i, time)$ for all $i = b, \dots, m$ - The most recent of all is preferred: if $T_j < T_k$ then $T_j \ll_{Cond} T_k$ **Conclusion -** $T_j \ll_{Cond} T_k$, then $G_P(n_j) \ll_{Time} G_P(n_k)$
Accuracy of derivation ($q_i = K3$, $Ptype = T2$) **Ptype -** $P_b = (G_M, C)$, $N_M = \{a\}$, $C : att(a, type) =$"PrimarySource", $\phi_{P_b}[G_P(i_j)] = \{a_a, \dots, a_l\}$, $\phi_{P_b}[G_P(i_k)] = \{a_b, \dots, a_m\}$ and shortest paths $Q_P(i_j) = \{S_{Pa}(i_j, a_a), \dots, S_{Pl}(i_j, a_l)\}$, $Q_P(i_k) = \{S_{Pb}(i_k, a_b), \dots, S_{Pm}(i_k, a_m)\}$ **Cond -** Given the longest paths $S_{Pj}(i_j, a_j)$ and $S_{Pk}(i_k, a_k)$: - $S_{Pj}(i_j, a_j) = (N_j, E_j)$, $\|E_j\| > \|E_i\|$, $E_i \subset S_{Pi}(i_i, a_i)$ for all $i = a, \dots, l$ - $S_{Pk}(i_k, a_k) = (N_k, E_k)$, $\|E_k\| > \|E_i\|$, $E_i \subset S_{Pi}(i_i, a_i)$ for all $i = b, \dots, m$ - The shortest path is preferred: if $\|E_j\| > \|E_k\|$ then $S_{Pj}(i_j, a_j) \ll_{Cond} S_{Pk}(i_k, a_k)$ **Conclusion -** $S_{Pj}(i_j, a_j) \ll_{Cond} S_{Pk}(i_k, a_k)$, then $G_P(n_j) \ll_{Der} G_P(n_k)$

Table 1. Procedure for establishing provenance criteria

6 A Formal Perspective

In this section we present a formalisation of the argumentation schemes Arg_A and Arg_B. We introduce the schemes in an argumentation framework for further investigation of the decision-making process. In particular, the framework must cater for defeasible rules and arguments about preferences, as this is required by the definition of our schemes. For this purpose, let us consider the Prakken and Sartor [13] approach, as it allows defeasible reasoning on preferences. Here, we only recall important aspects of the language.

The language of Prakken and Sartor [13] is formed by literals, input rules, and priorities. A literal is a ground atom L or its negation. Two types of negations are permitted: the classic negation $\neg$ and the negation as failure referred to as $\sim$.

A rule r is formed by literals connected by a two-place one-direction connective. A two-place predicate symbol $\prec$ is used to assert priorities between rules; e.g., $r \prec r'$. A *strong* literal is an atomic first-order formula, a formula of the form $r \prec r'$, or a formula preceded by strong negation $\neg$. The complement of an atom A is $\neg A$. In addition $\neg\neg A = A$. A *weak* literal is $\sim L$, where L is a strong literal. $\neg L$ represents the statement "L is definitely not the case", the weak literal $\sim L$ corresponds to "There is no evidence that L is the case".

Definition 9. *A* rule *is an expression of the form:*

$$r : L_0 \wedge \cdots \wedge L_j \;\; \wedge \sim L_k \wedge \cdots \wedge \sim L_m \Rightarrow L_n$$

where r indicates the name of the rule and each L_i is a strong literal.

The conjunction on the lefthand side of the arrow is referred to as the *antecedent*, and the righthand side one is the *consequent*. Two groups of input rules are identified: S the set of strict rules, and D the set of defeasible rules. Only defeasible rules can contain weak literals and in the following we will only consider defeasible rules, unless specified otherwise. The set of strict rules ensures a strict partial order of preferences; i.e., transitivity, contraposition of transitivity, and asymmetry.

Definition 10. *An argument is a finite sequence $A = [r_0, \ldots, r_n]$ of ground instances of rules such that:*

1. *for every i $(0 \leq i \leq n)$, for every strong literal L_j in the antecedent of r_i there is a $k < i$ such that L_j is the consequent of r_k;*
2. *No two distinct rules in the sequence have the same consequent.*

A literal L is a *conclusion* of A iff L is the consequent of some rule in A. A literal L is an *assumption* of A iff $\sim(\neg L)$ occurs in some rule in A. Let $A + S$ indicate the concatenation of an argument A and a set of strict rules S.

Definition 11. *Let A_1 and A_2 be two arguments, S_1, S_2 two sequences of strict rules, and $A_1 + S_1$ is an argument with conclusion L. Then*

1. *A_1 rebuts A_2 iff $A_2 + S_2$ is an argument with conclusion $\neg L$, provided that there is no rule in A_2 that is preferred than a rule in A_1.*
2. *A_1 undercuts A_2 iff A_2 is an argument with an assumption $\neg L$.*

Following [3], let us formalise Arg_A (Definition 1). In argument Arg_A, the 7W-statements of provenance warrant the assertion of a proposition i_j.

Definition 12 (Formalised argument Arg_A). *An instantiated argumentation scheme for provenance of type Arg_A is formed by:*

$$r_{A0} : \; \sim\neg evSupp_j \Rightarrow retrieved_j$$

$$r_{A1} : \; \bigwedge_{\ell=1}^{z} wj_\ell \wedge \sim\neg eps_j \wedge \sim\neg mpe_j \Rightarrow believable_j$$

$$r_{A2} : \; believable_j \wedge retrieved_j \Rightarrow holds_j$$

where the statements are:

- $holds_j$ = *"i_j holds"*
- $retrieved_j$ = "i_j was retrieved",
- $believable_j$ = "i_j is believable",
- $evSupp_j$ = "There is evidential support for i_j",
- eps_i = "Stated provenance elements are enough for believing i_j",
- mpe_j = "There are not missing provenance elements for i_j",
- wj_ℓ: *is a provenance element of* $G_P(i_j)$ *in the form of 7W-explanation,*
- $\wedge_{\ell=1}^{z} wj_\ell$: *is the stated 7W-explanation* ($z \leq 7$) *of a pattern* P_m *in* $G_P(i_j)$.

In our running example, consider $holds_j$ corresponding to state that i_j="The gang is heading south" holds. In Arg_A the initial P_g pattern of $G_P(i_j)$, shown in Figure 3, is represented with the 7W-statements, $\bigwedge_{\ell=1}^{4} w_\ell$ where:

- w_1 ="i_j was derived from Image of Border",
- w_2 ="i_j was associated with Reasoner",
- w_3 ="i_j was generated by the activity Image-processing",
- w_4 ="i_j was generated at time T_1".

The weak points of these rules are identified using the critical questions:

- Question **CQA1** stating "Is i_j consistent with other information?" points towards a rebuttal argument with conclusion $\neg holds_j$. In our example, a new argument Arg_A for $holds_k$ may represent an attack when $holds_k$ corresponds to state that i_k="The gang has just crossed the North border" holds. Furthermore, the analyst knows that if the gang is heading south, it could have not just crossed the northern border and vice-versa. This conflict is then represented by the following rules in the set of strict rules S:

$$r_{i0} : holds_j \Rightarrow \neg holds_k \qquad r_{i1} : holds_k \Rightarrow \neg holds_j \qquad r_{i0}, r_{i1} \in S$$

- Question **CQA2** states "Is i_j supported by evidence?". This undercuts the argument by challenging the assumption that the information about the gang heading south is based on some evidential support. An argument with conclusion $\neg evSupp_j$ will attack Arg_A.
- Questions **CQA3** and **CQA4** are directed to the assumptions in r_{A1}. CQA3 is "Does $G_P(n_i)$ contain other elements that lead us not to believe i_j?" and CQA4 is "Is there any other provenance element that should have been included to ensure that i_j is believable?". They refer to exceptions on the believability of the information based on provenance. CQA3 requires the agent to state more provenance elements. For example, we could add the pattern P_t about why i_j was gathered. CQA4 instead challenges the assumption that there are missing elements on the stated provenance pattern. For example, the image-processing that generates i_j did not use any previous observations. Arguments that attack Arg_A using CQA3 or CQA4 must contain a rule with conclusion $\neg eps_j$ or $\neg mpe_j$ respectively.

Similarly, argument scheme Arg_B (Definition 2) is formalised as:

Definition 13 (Formalised argument Arg_B). *An instantiated argumentation scheme for provenance preference of type Arg_B is formed by:*

$$
\begin{aligned}
r_{B0} &: \Rightarrow holds_j \\
r_{B1} &: \Rightarrow holds_k \\
r_{B2} &: \bigwedge_{\ell=1}^{u} wj_\ell \wedge \bigwedge_{\ell=1}^{v} wk_\ell \Rightarrow pref\_Ctr\_GP_{kj} \\
r_{B3} &: \sim\neg except\_Ctr\_GP_{kj} \wedge pref\_Ctr\_GP_{kj} \Rightarrow r_{B1} \succ r_{B0}
\end{aligned}
$$

where statements in the rules correspond to:

- $r_{B1} \succ r_{B0}$*: there is a preference* $i_j \ll i_k$ *between information* i_j *and* i_k*,*
- $pref\_Ctr\_GP_{kj}$ = "There is a criterion Ctr such that $G_P(i_j) \ll_{Ctr} G_P(i_k)$",
- $except\_Ctr\_GP_{kj}$ = "There is no exception to the criteria Ctr",
- wj_ℓ*: is a provenance element of* $G_P(i_j)$ *in the form of 7W-explanation,*
- wk_ℓ*: is a provenance element of* $G_P(i_k)$ *in the form of 7W-explanation,*
- $\wedge_{\ell=1}^{u} wj_\ell, \wedge_{\ell=1}^{v} wk_\ell$*: are the stated parts of provenance* ($u \leq |E_j|, v \leq |E_k|$)*.*

The critical questions point towards the construction of counterarguments:

- Question **CQB1** is "Does a different criterion Ctr_1, such that $\overline{G}_P(i_j) \gg_{Ctr_1} \overline{G}_P(i_k)$ lead $i_j \ll i_k$ not being valid?" Such a question can lead us to an argument rebutting Arg_B stating that there is another criteria Ctr_1 that asserts $G_P(i_j) \gg_{Ctr_1} G_P(i_k)$, with rule $pref\_Ctr1\_GP_{jk} \Rightarrow r_{B0} \succ r_{B1}$.
- Question **CQB2** states "Is there any exception to Ctr such that even if $G_P(i_j) \ll_{Ctr} G_P(i_k)$, $i_j \not\ll i_k$?" This identifies arguments with conclusion $\neg except\_Ctr\_GP_{kj}$. For example, this happens when one information item is more timely than another, but the time difference is irrelevant for the particular information.
- Question **CQB3** refers to "Is there any other reason for believing that $i_j \ll i_k$ is not valid?" This identifies an argument that has conclusion $\neg(r_{B1} \succ r_{B0})$. In this case the attacking argument provides reasons, unrelated to criteria for provenance, for which this preference is not valid.

Using the above arguments, agents will be able to reason about provenance elements of extracted information. Moreover, agents can rationally state preferences and resolve conflicting information.

6.1 Example of Decision-making

Let us develop further our running example where analyst Joe must decide if Gang G has crossed the northern border (see Figure 3). Here we describe how a software agent can support Joe in making this decision.

The information items are i_j= "The gang is heading south" and i_k="The gang has crossed the North border". Initially, our supporting agent states two arguments Arg_{A1} and Arg_{A2} for the generation paths P_g of i_j and i_k respectively.

Arg_{A1} states that i_j is credible because it was derived by the *Image of the Border*, and generated from *Image-processing* at time T_1 by the *Reasoner*. Arg_{A2} states that i_k is credible because the *Informer* has *Transcribed* the *Message Received* at time T_2. The conflict between Arg_{A1} and Arg_{A2} is described by CQA1.

The agent now constructs preferences according to different criteria. Assume that there are two preferences verified according to timeliness and trustworthiness. The agent has to be aware of Joe's order of trustworthiness attributed to actors; e.g., *"Reasoner"*$\ll$*"Informer"*. The preferences are presented to Joe with Arg_{B1} and Arg_{B2}. Arg_{B1} states that $i_j \gg i_k$ because in the asserted part of the provenance graphs (within Arg_{A1}-Arg_{A2}), the generation time of i_j, T_1, is more preferred than T_2, the generation time of i_k assuming $T_1 > T_2$. Arg_{B2} asserts the opposite, $i_j \ll i_k$ because the informer is considered more trustworthy than the reasoner. Hence, Arg_{B1} and Arg_{B2} rebut each other via CQB1.

In order to resolve these conflicts, we consider a preference between criteria stated by Joe (or the agent on his behalf). In the formalism, we include a rule that states a priority $\prec$ between the two priority rules r_{B3} within Arg_{B1} and Arg_{B2}. We assume that Joe asserts that *Trust* is preferred to *Time*, *Time* $\ll$ *Trust*. Then, "i_k holds" becomes the conclusion of an accepted argument. Joe receives from the agent the suggestion that the gang has crossed the northern border.

In our framework, we can also consider a conflict caused by a single criterion. When Joe asks for more information about provenance of i_j using CQA3, the agent reports an argument Arg_{A3} for the pattern P_g explaining how the message arrived to the informer. Arg_{A3} explains that the *Message Received* was *Delivered* by the *Messenger*. New preferences may now be asserted. For example, assume that Joe considers the messenger not to be trustworthy. Arg_{A3} changes the preference order of criterion *Trust* stated with Arg_{B2}. A new Arg_{B3} is stated for $i_k \ll i_j$ since the messenger is less trustworthy than the reasoner. The two rules of type r_{B2} within Arg_{B2} and Arg_{B3} have contradicting conclusions. A new rule is added to state the priority of the rule r_{B2} within Arg_{B3}, as this comprises a broader set of provenance elements. The agent informs Joe that the gang is heading south, since "i_j holds" is now the conclusion of an accepted argument.

We showed here how agents may support analysts in drawing conclusions exploring provenance. In future work, we will define a protocol for this interaction.

7 Discussion and Conclusion

In this paper, we discussed the importance for intelligence analysts of consulting provenance information for identifying plausible explanations for a situation. We have presented a model of argumentation schemes that contributes towards the exploration and the assessment of provenance in a structured format. We proposed an argument scheme for presenting provenance data that facilitates the assessment of whether a piece of information is credible. A second argument scheme states preferences between pieces of information based upon provenance criteria for intelligence analysis. These schemes select information that comes from more reliable, trustworthy, accurate, or more timely retrieval processes.

We envisage this method to contribute in the broader process of intelligence analysis permitting the analysts to focus on more credible claims.

Provenance is a novel application for argumentation-based frameworks. The approach that first discussed the use of arguments underpinned by provenance is Chorley et al. [2], where provenance is recorded for justifications provided by users during the assessment of policy options. An analyst identifies more suitable policies by reasoning about justifications and their provenance. In contrast, our goal here is to transfer this workload to agents by introducing provenance within the automated reasoning process. In the context of trust, the work of Parsons et al. [9], analyses the sources of information to infer the trustworthiness of arguments. However, there are other elements of provenance, in particular related to the manipulation of information, that we consider in this research.

Using provenance for assessment of information quality has also been explored. Hartig and Zhao [5] proposed a measure of timeliness using a specific model of provenance, according to creation and access time. In our model, we provide simple, intuitive criteria to represent how intelligence analysts would evaluate information based on provenance elements. In future work, we will investigate more complex quality measures such as those proposed in [10]. Our criteria, however, permit agents to instantiate concrete preferences providing a domain of application for argumentation frameworks such as [7, 13].

The argumentation framework discussed in this paper is based on the Prakken and Sartor [13] approach, as it provides useful insights on the decision-making process with defeasible preferences. In future work, we will consider other frameworks (e.g. [12]) that may provide different characteristics for further developments of our work. Furthermore, we have only considered a small set of criteria. Another criterion may, for example, be based on the fact that an activity has some input requirements that if not satisfied result in an incorrect output even if the activity is reliable. Another may assess the complexity of computing the derivation chain of an information item; e.g., a long chain of simple arithmetic operations could be more accurate than a single step chain that involves the result of an optimisation algorithm. We will investigate interactions between criteria; e.g., the relationship between trustworthy actors and reliable activities. We will also discuss the possibility for agents to adopt subjective preferences. Ultimately, our aim is to evaluate and enrich our approach with the help of domain experts [14].

In this research, we addressed the problem of facilitating analysts in evaluating hypotheses by exploring and assessing the provenance of supporting evidence. We demonstrated through examples, that agents employing our model of argumentation schemes will be able to support analysts in assessing the credibility of information and in resolving conflicts by reasoning about provenance.

Acknowledgments. The authors would like to thank Edoardo Pignotti at the University of Aberdeen for the useful discussion on provenance. This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911 NF-06-3-

0001. The views and conclusions contained in this document are those of the author(s) 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.

References

1. Bex, F., Prakken, H., Reed, C., Walton, D.: Towards a formal account of reasoning about evidence: argumentation schemes and generalisations. Artificial Intelligence and Law 11(2-3), 125–165 (2003)
2. Chorley, A., Edwards, P., Hielkema, F., Philip, L., Farrington, J.: Supporting provenance and argumentation in evidence-based policy assessment. In: Proceedings of the Oxford eResearch Conference (2008)
3. Croitoru, M., Fortin, J., Oren, N.: Arguing with preferences in EcoBioCap. In: Proceedings of the Fourth International Conference on Computational Models of Argument. pp. 51–58 (2012)
4. Goble, C.: Position statement: Musings on provenance, workflow and (semantic web) annotations for bioinformatics. In: Workshop on Data Derivation and Provenance, Chicago (2002)
5. Hartig, O., Zhao, J.: Using web data provenance for quality assessment. Proceedings of the First International Workshop on the Role of Semantic Web in Provenance Management (2009)
6. Lim, C., Lu, S., Chebotko, A., Fotouhi, F., Kashlev, A.: OPQL: querying scientific workflow provenance at the graph level. Data & Knowledge Engineering 88(0), 37 – 59 (2013)
7. Modgil, S.: Reasoning about preferences in argumentation frameworks. Artificial Intelligence 173(9–10), 901 – 934 (2009)
8. Moreau, L., Missier, P.: PROV-DM: The PROV Data Model (March 2013), available at http://www.w3.org/TR/prov-dm/
9. Parsons, S., Tang, Y., Sklar, E., McBurney, P., Cai, K.: Argumentation-based reasoning in agents with varying degrees of trust. In: Proceedings of the Tenth International Conference on Autonomous Agents and Multiagent Systems. vol. 2, pp. 879–886 (2011)
10. Pipino, L.L., Lee, Y.W., Wang, R.Y.: Data quality assessment. Communications of the ACM 45(4), 211–218 (2002)
11. Pirolli, P., Card, S.: The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In: Proceedings of the International Conference on Intelligence Analysis (2005)
12. Prakken, H.: An abstract framework for argumentation with structured arguments. Argument and Computation 1(2), 93–124 (2010)
13. Prakken, H., Sartor, G.: Argument-based extended logic programming with defeasible priorities. Journal of Applied Non-Classical Logics 7(1-2), 25–75 (1997)
14. Various Intelligence Analysts: Personal communication. NATO analyst exchange meeting (September, 2013)
15. Walton, D., Reed, C., Macagno, F.: Argumentation schemes. Cambridge University Press (2008)