



City Research Online

City, University of London Institutional Repository

Citation: Arunkumar, S., Sensoy, M., Srivatsa, M. & Rajarajan, M. (2015). Reasoning with Streamed Uncertain Information from Unreliable Sources. *Expert Systems with Applications*, 42(22), pp. 8381-8392. doi: 10.1016/j.eswa.2015.04.031

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/8678/>

Link to published version: <https://doi.org/10.1016/j.eswa.2015.04.031>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Reasoning with Streamed Uncertain Information from Unreliable Sources

Saritha Arunkumar^{a,*}, Murat Sensoy^d, Mudhakar Srivatsa^d, Muttukrishnan Rajarajan^a,

^aIBM Hursley Labs, UK

^bUniversity of Aberdeen, UK

^cIBM Research, USA

^dCity University London, UK

Abstract

Humans or intelligent software agents are increasingly faced with the challenge of making decisions based on large volumes of streaming information from diverse sources. Decision makers must process the observed information by inferring additional information, estimating its reliability, and orienting it for decision making. Processing streaming trust framework, when fact is getting created and inferred is a process in online mode and our paper works efficiently in online mode. In offline mode, someone initiates a query and gets an output based on the query. In this paper we have mainly shown that unstructured reports from unreliable and heterogeneous sources are processed to generate structured information in Controlled English. Uncertainty in the information is modelled using Subjective Logic that allows statistical inference over uncertain information. Trustworthiness of information is modelled and conflicts are resolved before fusion. This process is totally undertaken on streaming information resulting in new facts being inferred from incoming information which immediately goes through trust assessment framework and trust is propagated to the inferred fact. In this paper, we propose a comprehensive framework where unstructured reports are streamed from heterogeneous and potentially untrustworthy information sources. These reports are processed to extract valuable uncertain information, which is represented using *Controlled Natural Language* and *Subjective Logic*. Additional information is inferred using deduction and abduction operations over subjective opinions derived from the reports. Before fusing extracted and inferred opinions, the framework estimates trustworthiness of these opinions, detects conflicts between them, and resolve these conflicts by analysing evidence about the reliability of their sources. Lastly, we describe an implementation of the framework using International Technology Alliance (ITA) assets (Information Fabric Services and Controlled English Fact Store) and present an experimental evaluation that quantifies the efficiency with respect to accuracy and overhead of the proposed framework.

Keywords: information fusion, controlled natural language, subjective logic

*Corresponding author

Email addresses: saritha.arun@uk.ibm.com (Saritha Arunkumar), m.sensoy@abdn.ac.uk (Murat Sensoy), msrivats@us.ibm.com (Mudhakar Srivatsa), R.Muttukrishnan@city.ac.uk (Muttukrishnan Rajarajan)

Preprint submitted to Elsevier

April 7, 2015

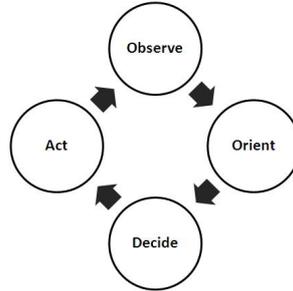


Figure 1: The OODA loop

1. Introduction

A number of effective decision-making processes devised to support strategic operations and/or tactical missions and tasks in both military and civilian settings follow a recurring pattern from information gathering to action taking. This pattern is summarily represented by USAF Col. John Boyd's OODA loop (Brehmer, 2005) containing the following four phases (see Figure 1):

- The *Observe* phase comprises the processes of accumulating information pertinent to a situation of interest.
- The *Orient* phase comprises the processes of analyzing the information within the context of the particular task (operation or mission) at hand.
- The *Decide* phase comprises the processes of determining a proper course of action based on the knowledge gained during the *Orient* phase.
- The *Act* phase comprises the processes involving the actual (physical) actions taken.

The OODA phases continuously interact with each other, with, for example, the *Orient* phase *requesting* additional information from the *Observe* phase to drive better decisions. The interpretation of facts and information analysis in the *Orient* phase is heavily influenced by the existing (background) knowledge and understanding of the current situations at various levels including social, political, economic, cultural, administrative, logistic capabilities, resource availability, past experiences, and so on. In proposing the OODA loop, Col. Boyd had argued that in conflict situations, such as when acting to intercept an enemy airplane, in order to prevail one must execute his own OODA cycle at a pace faster than that of his opponent's OODA loop causing the opponent confusion by acting on stale information from situations that have already changed.

Our interest in this work are in the first two phases that deal with the gathering, processing, and contextual placing of information to drive subsequent decisions and actions. It is motivated by the fact that decision makers (humans or software agents alike) are increasingly faced with the challenge of examining large volumes of information originating from heterogeneous sources requiring to ascertain opinion in various pieces of arriving information. Detection of conflicts between opinions and resolution of conflicts is a very important part of the solution. Resolving conflicts by analysing evidence about the reliability of the sources is shown very well in this paper. The challenge is exacerbated when the decision makers should employ uncertain rules

to reason about the information collected from these heterogeneous sources. Unfortunately, this is the case in most settings, since inferencing rules are uncertain and vague in many sensing domains and scenarios.

In typical applications of the OODA loop, the information sources feeding the *Observe* and supporting the *Decide* phases, such as the intelligence gathering scouts or the visual and electromagnetic sensors on a fighter airplane, are assumed to be administratively associated (if not physically owned and controlled) by the OODA beneficiary, e.g., the mission planners or the fighter pilot. In such cases, the capabilities of these sources are typically well known and documented and this background knowledge can be accounted for sufficiently well during the *Orient* phase. However, in coalition environments involving loosely coupled sensory systems (Sensoy et al., 2012) ownership, control and effectiveness of the sources cannot be guaranteed. The sensing resources may belong to different coalition partners or the informants may be locals. In such cases, the provenance and ultimately the *quality of the information* (Bisdikian et al., 2009, 2011) derived from these sources becomes questionable. Therefore, to properly support decision and action processes, the *Observe* and *Orient* phases must explicitly deal with uncertain information from sources whose capabilities are only partially (if at all) known and documented. In other words, the level of *trustworthiness* of the gathered information needs to be accounted for and properly weighed in the decisions and actions taken. There are various definitions of trust. McKnight and Chervany's work (Jøsang and Presti, 2004) defines trust as follows: Trust is the extent to which one party is willing to depend on somebody, or something, in a given situation with a feeling of relative security, even though negative consequences are possible. We follow this definition as it is used in most computer science publications and applications.

It is thus an objective of this work to design an opinion assessment framework that "scores" information arriving from uncertain sources and resolved conflicts. In supporting the *Observe* and *Orient* phases of the OODA loop, the framework will accommodate streaming information, acknowledging the need to deal with fast-paced information, such as that derived from sensors deployed across geographical spreads.

The main significance of this work shows the information being processed from streaming data, facts derived from it and new facts inferred from it. When new facts are inferred, trust is assessed for the new facts. So the process is executed on a continuous online mode. This is main difference between the previous works in this area and our work. In the offline mode, the facts are not processed immediately and the trust is not assessed instantaneously.

Another key observation from this paper is that the work introduces key facts being derived from streaming data and the key facts lead to new facts being inferred. As the new facts are inferred it is passed through the trust assessment framework to propagate a trust value to the new inferred fact.

This solution is very applicable in social computing and social networks (Wang et al., 2015) Social networking applications like twitter can be used to collect information, our information sources are users, information is received as unformatted text which needs to be converted into Controlled English format. The formatted text would infer new facts and a trust assessment framework is used to assess the trust of the new fact.

The organization of the paper is as follows: Section 2 presents background and highlights features of our assessment framework along with an exemplar use case. Section 3 provides a brief introduction to *subjective logic* that we use in our framework to assess opinions in Section 4. Section 4 describes the opinion assessment framework and its components. Section 5 describes how conflicts between opinions are detected and resolved before during information fusion. Section 6 describes an implementation of our framework on the Information Fabric (Wright et al., 2009)

and evaluates the performance and accuracy of our opinion assessments. Section 7 discusses the related work and finally, Section 8 concludes the paper with an overview of contributions and directions for future work.

2. Background

As information accumulates and is read for consumption, its trustworthiness must be assessed for weighing properly in subsequent decision-making. Several authors have explored various opinion computation models on static data (e.g., eBay recommendation system (Schafer et al., 1999), Netflix movie ratings, EigenTrust (Kamvar et al., 2003), PeerTrust (Xiong and Liu, 2004), etc.). A common data model subsumed by several trust computation models (as succinctly captured in Kuter and Golbeck (Kuter and Golbeck, 2007)) is the assignment of numeric trust scores between pairs of entities (e.g., in eBay recommendation buyers rate sellers, in Netflix movie ratings users rate movies, etc.). Such pair-wise numeric ratings contribute to a (dis)similarity score (e.g., based on the \mathcal{L}_1 or \mathcal{L}_2 norms, cosine distance, etc.) which is used to compute personalized trust scores (as in PeerTrust) or recursively propagated to compute global trust scores (as in EigenTrust).

There are two common assumptions in several such opinion models that are viable in commercial settings: (i) a statistically significant number of trust evidence (e.g., ratings) is available prior to opinion assessment, and (ii) assessed opinion values tend to vary slowly over time. In contrast, military settings warrant: (i) trust assessment over partial, uncertain and streaming (live and real-time) information from heterogeneous sources, and (ii) coping up with the dynamic and evolving nature of the ground truth. While most of the past work is focused on static data, the goal of this work is to develop a family of opinion operators for such dynamic information flows, i.e., assess opinion over *data-in-motion* rather than a large corpus of static data.

2.1. Use Case and System Highlights

Let us consider the following scenario comprising the events:

- *Event e_0* : An explosion occurs at location l_0 at time t_0 . Event e_0 represents the ground truth and is assumed that there is no direct evidence available to support it in the opinion assessment system.
- *Event e_1* : An acoustic sensor reports a possible explosion at location l_1 at time t_1 .
- *Event e_2* : A human agent reports unusual dust level at location l_2 at time t_2 .

In this scenario, the goals are two-fold: (i) fact inference (e.g., infer an explosion from event e_2) and information fusion (e.g., fuse events e_1 and e_2 based on their spatio-temporal proximity), and (ii) assessing additional information (metadata) about the inferred/fused facts, such as an estimate of where and when (\hat{l}_0 and \hat{t}_0) the explosion took place, or an assessment of opinion in the fact that an explosion occurred based on previously seen events.

With reference to Figure 2, the first goal is addressed by the analytics component, the design of which is outside the scope of this paper. In general, we expect to deal with a live (and potentially bursty) stream of events (e.g., textual or sensory reports) that are tagged with metadata (e.g., location, time, owner/author-ship, beliefs). The stream of events may be subject to various analytics operations, e.g., fact extraction from event data using the Controlled English (Mott, 2010) Store, event data fusion using the Information Fabric (Wright et al., 2009), etc. We do

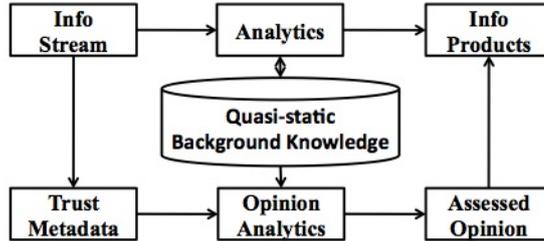


Figure 2: Opinion Assessment on Streaming Information

not make any specific assumptions on the nature of such analytics operations, but for the fact that they are constrained to operate on streaming data. Hence, given latency constraints, storage, archival and post processing (or even batch processing) of event data may not be a viable option. We remark that this does not preclude the option of storing compact summaries (e.g., in a CE fact store (Mott, 2010)); only storing all (or a significant portion of) historical events is deemed infeasible. We also note that the goal of this paper is not to support such analytics operations; instead our goal is to augment such operations with opinion metadata assessment.

While the internal workings of the analytics component are isolated from the opinion assessment system, we assume that the background knowledge is shared (e.g., rules for Bayesian inference, Controlled English fact inference rules, etc.). Intuitively, we model the analytics operation as accepting one or more inputs and a quasi-static background knowledge and generating one or more outputs. Here, the “quasi-static” qualifier means that the background knowledge is updated at a rate that is significantly slower than the event rate. This background knowledge is sometimes also referred to as the *domain knowledge*.

We model opinion assessment as a family of operators that accepts input opinion metadata and background knowledge and computes output opinion metadata. This model only requires that the background knowledge be shared between the analytics and the opinion assessment. Typically, such background knowledge is compact and has well defined semantics and isolates the opinion assessment system from understanding all the intricate details of the analytics operations (see Figure 2). Architecturally speaking, opinion is assessed as a part of a call-back function that is invoked each time the analytics operation generates a new output; the call-back function is supplied with the opinion metadata of the relevant inputs that were used to derive the said output.

Another interesting component of this work is the representation of rules with opinions. Past work assumed that the rules were always certain. Through this paper, we have experimented the rules with opinions and certain uncertainty using the Subjective logic triples, i.e., *belief*, *disbelief*, and *uncertainty*.

In this paper we present a realization of this model using subjective logic triples (i.e., belief, disbelief and uncertainty) for representing opinion and Bayesian inference network for the background knowledge. In this setting, we have built a family of subjective logic operators that can be used to compute opinion metadata for an output event, given the opinion metadata on the input events and a Bayesian network that encodes probabilistic dependencies between one or more facts, e.g., $\Pr(\text{dust}|\text{explosion}) = 1$, $\Pr(\text{dust}|\neg\text{explosion}) = 0.25$, and these probabilities may change over a time horizon that is significantly longer than the duration of an event burst. Later in the paper, we also present an early implementation of our opinion assessment framework within

the Information Fabric along with some experimentation results from the implementation.

In closing this section, we note that in Figure 2, we have separated the functional modules of opinion assessment in two parts: (i) the part that represents the incoming information product (Info Stream) and the processes that “blindly” operate on it, such as for extracting metadata (opinion Metadata), and (ii) the part that qualifies the information or its metadata, makes inferences using it, provides context and produces outgoing information products. While the OODA loop does not represent a system architecture, it, nonetheless, would be natural to associate the modules in the first part with the *Observe* phase as they relate with the accumulation of information, and the modules in the second part with the *Orient* phase as they relate with providing context and assessments within that context for the information.

3. Subjective Logic

Subjective logic (SL) is a type of probabilistic logic that explicitly takes uncertainty and belief ownership into account. In general, SL is suitable for modelling and analysing situations involving uncertainty and incomplete knowledge (Jøsang, 2011). Arguments in SL are subjective opinions about propositions.

A binomial opinion of an agent A about the truth of a proposition x is represented by the quadruple $w_x^A = (b, d, u, a)$, where: b is the belief that x is true; d is the belief that x is false; u is the uncertainty (or, uncommitted belief) about x ; and a is the base rate about x which represents the a priori probability about x in the absence of evidence; and $b + d + u = 1.0$ and $b, d, u, a \in [0, 1]$. Characteristics of a binary subjective opinion can be summarised as follows:

- $b = 1$ then it is equivalent to binary logic *TRUE*
- $d = 1$ then it is equivalent to binary logic *FALSE*
- $b + d = 1$ then it is equivalent to traditional probability
- $b + d < 1$ then it expresses some degrees of uncertainty
- $b + d = 0$ then it expresses total uncertainty

A binomial opinion can be represented as a Beta distribution. The probability expectation value of an opinion is defined as $E(w_x^A) = b + u \times a$, based on the corresponding Beta distribution (Jøsang, 2011). Through the correspondence between binomial opinions and Beta distributions, SL provides an algebra for subjective opinions. Some of the operators defined over subjective opinions are listed in Table 1. Detailed descriptions and proofs for soundness of these operators are out of the scope of this paper, but they can be found elsewhere (Jøsang, 2011, 2001, 2002; Jøsang and Grandison, 2003; Jøsang and McAnally, 2005; Jøsang, 2007).

Each opinion $w_x = (b, d, u, a)$ for a proposition x implies another opinion $w_{\neg x} = \neg w_x = (d, b, u, 1 - a)$. For instance, let x be the proposition “Dust storm in location X ”, the a priori probability for x be $\frac{1}{2}$, and the opinion of sensor s_1 about x be $w_x^{s_1} = (0.75, 0.1, 0.15, \frac{1}{2})$, then the sensor’s opinion about “No dust storm in location X ” would be $w_{\neg x}^{s_1} = (0.1, 0.75, 0.15, \frac{1}{2})$. If another sensor s_2 has another opinion $w_x^{s_2} = (0, 0.7, 0.3, \frac{1}{2})$, these two opinions are fused into a single opinion using the consensus (a.k.a., cumulative fusion) operator \oplus (Jøsang, 2007): $(0.75, 0.1, 0.15, \frac{1}{2}) \oplus (0, 0.7, 0.3, \frac{1}{2}) = (0.55, 0.34, 0.11, \frac{1}{2})$. The consensus operator is defined as follows, see (Jøsang, 2002).

Table 1: Some operators for binary opinions in Subjective Logic (Jøsang, 2011).

Operator Name	Notation
Complement (Jøsang, 2001)	$w_{\neg x} = \neg w_x$ $\neg(b, d, u, a) = (d, b, u, 1 - a)$
Discounting (Jøsang, 2001)	$w_x^{A:B} = w_B^A \otimes w_x^B$
Fusion (Jøsang, 2002, 2007)	$w_x^{A \diamond B} = w_x^A \oplus w_x^B$ $w_x^{1 \diamond \dots \diamond n} = \oplus \{w_x^1, \dots, w_x^n\}$
Disjunction / OR (Jøsang and McAnally, 2005)	$w_{x_1 \vee x_2} = w_{x_1} \sqcup w_{x_2}$ $w_{x_1 \vee \dots \vee x_m} = \sqcup \{w_{x_1}, \dots, w_{x_m}\}$
Conjunction / AND (Jøsang and McAnally, 2005)	$w_{x \wedge y} = w_x \cdot w_y$ $w_{x_1 \wedge \dots \wedge x_m} = \cdot \{w_{x_1}, \dots, w_{x_m}\}$
Deduction (Jøsang and Grandison, 2003)	$w_{y x} = w_x \odot (w_{y x}, w_{y \neg x})$
Abduction (Jøsang and Grandison, 2003)	$w_{y x} = w_x \bar{\odot} (w_{y x}, w_{y \neg x})$

Definition 1. Let $w_x^A = (b_x^A, d_x^A, u_x^A, a_x^A)$ and $w_x^B = (b_x^B, d_x^B, u_x^B, a_x^B)$ be opinions respectively held by agents A and B about the proposition x. When $u_x^A, u_x^B \rightarrow 0$, the relative dogmatism between w_x^A and w_x^B is defined by $\gamma = u_x^B / u_x^A$. The fused opinion $w_x^{A \diamond B} = (b_x^{A \diamond B}, d_x^{A \diamond B}, u_x^{A \diamond B}, a_x^{A \diamond B})$ can be computed as follows, where $\kappa = u_x^A + u_x^B - u_x^A u_x^B$:

for $\kappa \neq 0$	for $\kappa = 0$
1. $b_x^{A \diamond B} = (b_x^A u_x^B + b_x^B u_x^A) / \kappa$	1. $b_x^{A \diamond B} = (\gamma b_x^A + b_x^B) / (\gamma + 1)$
2. $d_x^{A \diamond B} = (d_x^A u_x^B + d_x^B u_x^A) / \kappa$	2. $d_x^{A \diamond B} = (\gamma d_x^A + d_x^B) / (\gamma + 1)$
3. $u_x^{A \diamond B} = (u_x^A u_x^B) / \kappa$	3. $u_x^{A \diamond B} = 0$
4. $a_x^{A \diamond B} = (a_x^A u_x^B + a_x^B u_x^A - (a_x^A + a_x^B) u_x^A u_x^B) / \kappa$	4. $a_x^{A \diamond B} = (\gamma a_x^A + a_x^B) / (\gamma + 1)$

Discounting operator \otimes allows normalisation of opinions based on the trustworthiness of their owners (Jøsang, 2001). In order to allow different trust models in our work easily, we adopt a slightly modified version of the discounting operator. Let us assume that we get the opinion (b_x^i, d_x^i, u_x^i, a) from source i about the fact x and have t_i as the trustworthiness of the source. The resulting opinion normalised based on opinion is computed as

$$(b_x^i \times t_i, d_x^i \times t_i, u_x^i + (1 - t_i) \times (b_x^i + d_x^i), a)$$

where the idea is to move a percentage of belief and disbelief to uncertainty based on the level of trust in the source. The trustworthiness of information sources can be calculated using different models. For example, beta reputation system (Jøsang and Ismail, 2002) uses beta probability density functions (Jøsang and Ismail, 2002) to model trust. A Beta distribution has two parameters $(r_i + 1, s_i + 1)$, where r_i is the amount of positive evidence and s_i is the amount of negative evidence for the trustworthiness of the source i . The degree of trust t_i is then computed as the expectation value of the Beta distribution:

$$t_i^j = \frac{r_i^j + 1}{r_i^j + s_i^j + 2} \quad (1)$$

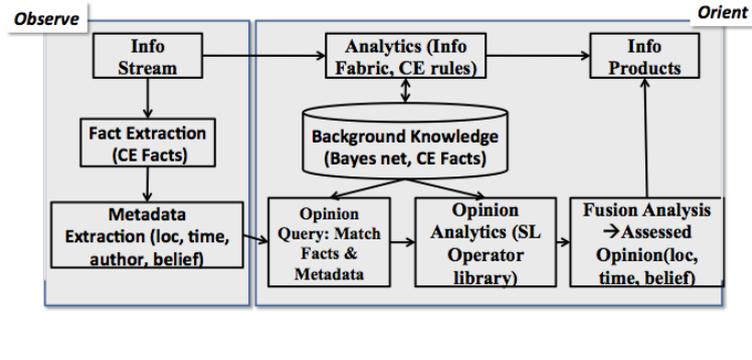


Figure 3: Opinion Assessment Framework

While this computation is simple and straight forward, there exist more complicated approaches for modelling trust. Our work is flexible enough to integrate any of such model as long as these models provide a trust value in range $[0, 1]$.

Now, we give a concrete example of discounting operation. Let the trustworthiness of s_1 and s_2 be 0.1 and 0.8, respectively. Then, normalised opinions of s_1 and s_2 would be $w_x^{s_1} \otimes 0.1 = (0.07, 0.01, 0.92, \frac{1}{2})$ and $w_x^{s_2} \otimes 0.8 = (0.0, 0.56, 0.44, \frac{1}{2})$, respectively. In this case, fusion of the normalised opinions would be $(0.03, 0.54, 0.43, \frac{1}{2})$. In the rest of the paper, and unless otherwise specified, we will assume that *all opinions are normalised based on trustworthiness of their owners* (i.e., sources).

Given opinions about different propositions (e.g., w_{x_1} and w_{x_2}), SL provides conjunction (\cdot) and disjunction (\sqcup) operators to compute opinions for the conjunction (e.g., $w_{x \wedge y}$) and disjunction (e.g., $w_{x \vee x_2}$) of the propositions (Jøsang and McAnally, 2005). In addition to these operators, deduction operator (\odot) is used to deduce new opinions from existing ones (Jøsang and Grandison, 2003). For instance, let x be the proposition “Country A is buying missiles” and y be “Country A is preparing to attack”. Let us assume that we have the following conditional opinions about “ y given x ” and “ y given $\neg x$ ”.

- $w_{y|x} = (0.9, 0.0, 0.1, \frac{1}{2})$
- $w_{y|\neg x} = (0.0, 0.0, 1.0, \frac{1}{2})$

Given an opinion about x , we can deduce an opinion for y using $w_{y||x} = w_x \odot (w_{y|x}, w_{y|\neg x})$ as described in (Jøsang and Grandison, 2003). If the opinion about x is $w_x = (0.75, 0.25, 0.0, \frac{1}{2})$, the deduced opinion about y is computed as $w_{y||x} = (0.68, 0, 0.32, \frac{1}{2})$ based on the algorithm in (Jøsang and Grandison, 2003). Details of the deduction operator is not in the scope of this paper, but it can be found in (Jøsang and Grandison, 2003).

4. Opinion Assessment Framework

Figure 3 shows the architecture of our opinion assessment framework; wherever appropriate, we indicate ITA assets that could be leveraged in implementing the framework. As described earlier, we assume information analytics operates over a stream of input information and generates

information products e.g., by fusing information, by extracting facts and using inferences to derive new facts, etc. Similarly to Figure 2, in Figure 3 we also show those parts of the architecture that relate to the *Observe* and *Orient* phases of the OODA loop.

Our opinion assessment framework operates in parallel with the information processing pipeline. The opinion assessment framework extracts facts and opinion metadata from the stream of input information. The background knowledge store is queried for identifying relevant facts, i.e., facts whose opinion level may influence or may be influenced by the opinion level of the input information. For example, the input information may indicate unusual dust levels and background knowledge may encode the fact that unusual dust levels are often a result of an explosion; hence, the fact store may be queried for potential explosion information. Also, the background knowledge may encode the fact that explosion may be identified using its acoustic signature; hence, the fact store may be queried for acoustic information from a relevant space-time region (volume, to be more accurate) that surrounds the unusual dust level.

Once relevant facts are extracted, opinion in input information and information products is assessed using the family of opinion operators built using subjective logic. The choice of the opinion operator depends upon the nature of analytics operation (e.g., conjunction vs. deduction vs. consensus). In this section we present three examples which illustrate the functionality of various components in our opinion assessment framework.

One of our key design features is that of decoupling the analytics component with the opinion assessment system. The analytics component is responsible for extracting and inferring facts (e.g., from unstructured text or other primitive facts). In doing so the analytics component may leverage (e.g., use a static road/terrain map to enrich facts) and update (e.g., Controlled English (CE) (Mott, 2010) rules to infer new facts) background knowledge. The opinion assessment system, on the other hand, does not infer new facts (as indicated by the unidirectional flow of information from background knowledge to the opinion assessment system in the figure). Instead it augments the inference task in the analytics component with opinion assessment capability, i.e., assesses opinion on the inferred fact.

In the rest of this section, we describe each of these steps in detail.

4.1. Fact Extraction

The input to our system is typically a stream of unstructured reports. Each report is subjected to fact extraction, resulting in a set of facts that are subsequently used by both the analytics and the opinion assessment components. The extracted facts and texts may be stored in Controlled English (CE) format in a CE Store. Here is a snippet of the intelligence report.

For example, given an intelligence report the following facts may be extracted¹:

- there is a group named ent1_wisrep08 that has ‘a foot patrol’ as description
- there is an entity reference named er1_wisrep08 that has ‘a foot patrol’ as description and
 - has the document wisrep08 as source and
 - has the group ent1_wisrep08 as entity and
 - has 40 as start position and
 - has 53 as end position

¹We use teletype lettering to emphasize CE specific word patterns.

WISREP	
Report number:	WIS/AF/008
Report date:	14th February 2010
Event date and time:	Disabling a potential suicide bomber wearing a vest
Event location:	Khoshal Village, Nad-e Ali District
Target	
	The target appears to have been a foot patrol operating in KHOSHAL village. The suicide bomber, a young man with an age thought to be around 15-16 years, was wearing the vest and had been sent from a terrorist stronghold to attack the foot patrol when a dicking network had reported the patrol had entered the village market area. He was riding a motorcycle which was also loaded with explosives disguised as clothing.
Device 1	

Figure 4: A Snippet of an Unstructured Report

4.2. Metadata Extraction

Beyond fact extraction, we also extract various metadata attributes from a report including, Creator, Keywords, Opinion, Subject, Title, Security Level, Company, Category, Document parts, etc. For example,

- Source = “Dr Dave Sloggett”
- has ‘Opinion’ as (0.7, 0.2, 0.1, 0.5)
- Keywords = “IED, Pathfinder”
- Subject = “Daisy Chain Device”
- Title = “WISREP 001”
- Security Level = 0
- msol:codepage = 1252
- Company = “someCompanyName”
- Category = “WISREP”
- Document Parts = [(0, “WISREP 001”)]

Metadata attributes thus generated are also asserted as facts, thereby, permitting a uniform representation for both data and metadata. For example, Document ‘20110228_WISREP_Report_001_ISAF’

- has the ms-word metadata ‘20110228_WISREP_-Report_001_ISAF.doc.meta.txt’ as metadata and
- has ‘ISAF’ as filename security level and

- has ‘001’ as filename subtitle and
- has ‘WISREP’ as filename report type and
- has ‘20110228’ as filename date and
- has ‘Report’ as filename title and
- has ‘20110228_WISREP_Report_001_ISAF.doc.plain.txt’ as plain filename and
- has ‘Wed Jun 8 13:24:33 2011’ as time processed and
- has ‘PFIII document processor’ as agent and

It is important to emphasize the following about the reports, facts, and opinions. Within the same report, there can be multiple facts written by multiple sources. When a fact is extracted from a report, the meta-data includes the source and opinion asserted by the source about this fact. Before asserting the opinion from the specific source into the knowledge base, we discount it with the trustworthiness of the source. If the opinion is not asserted by the source in the report, we take it as (1, 0, 0) by default and discounted based on opinion. For instance, if the trustworthiness of the source is (0.5, 0.1, 0.4), the opinion about the fact is taken as (0.5, 0, 0.5) after discounting. Normally, the fact may appear in different reports. Therefore, there may be many opinions about the facts in our knowledge base. Some of these facts may be in conflict. For instance, we may derive opinions (0.8, 0, 0.2) and (0.1, 0.8, 0.1) for the same fact “there is a problem in the road”. The first opinion indicates that the fact is probably true while other implies that the fact is not true. We will later describe how to handle these conflicting opinions about the same facts.

4.3. *Opinion Query*

Each opinion is associated with a proposition, which may be a fact or meta-data. For example, the knowledge base stores the opinion (0.9, 0, 0.1) from USGS about the fact “seismic activity has happened at location X on 20th June 2014 at time 16:01” and the opinion (0.8, 0, 0.2) for the proposition “USGS is a reliable source for seismic data”. In this section, given a specific fact extracted from a report, we describe how to identify relevant facts (and meta-data), and retrieve opinions related to these. In particular, we extract three types of relevant facts and retrieve associated opinions from the knowledge base:

- The facts that exhibit causal relationships with the input fact (e.g., dust level, acoustic sensor measurements with an IED explosion)
- The facts that exhibit spatio-temporal similarity (e.g., two explosion reports whose space-time tags are relatively close to each other)
- The facts with matching meta-data attributes such as creator, keywords, etc.

After retrieving all the known opinions about the relevant facts and meta-data, we perform opinion analysis as described in the next section.

4.4. Opinion Analytics

The opinion analytics component is designed to operate in parallel with the main analytics component; one could view our opinion analytics component as being triggered by the main analytics component whenever new facts are inferred or updated by the analytics component. The trigger would include as references the input and output facts from the analytics component and the relevant background knowledge queried in the previous step (Opinion Query step). The opinion analytics component uses a subjective logic (SL) operator library for deducing new opinions from the existing opinions about known facts. In this section we describe three sample opinion analytics operations.

4.4.1. Forward Chaining (input opinions are known)

Consider two known facts that are annotated with metadata such as source (e.g., coalition member) C , location L and time T and an opinion tv . Given two annotated facts $(f_1, c_1, l_1, t_1, tv_1)$ and $(f_2, c_2, l_2, t_2, tv_2)$, the analytics operation may infer and generate a new opinion about another fact f_3 using deduction as: $f_1 \text{ AND } f_2 \Rightarrow f_3$. The goal of the opinion assessment system in this case is to infer the opinion in the fact f_3 . In order to do this, the opinion analytics system would identify facts f_1 and f_2 and the inference rule used to generate fact f_3 , namely, $f_1 \text{ AND } f_2 \Rightarrow f_3$. In this example, the conjunction of the opinions about f_1 and f_2 are computed using the conjunction operator of SL, which is described in Definition 1. The resulting opinion is the opinion for the left hand side of the rule, i.e., $f_1 \text{ AND } f_2$. Based on the computed opinion and the deduction operator \odot mentioned in Section 3, we compute the opinion tv_3 associated with the fact f_3 .

Definition 1. Let $w_x = (b_x, d_x, u_x)$ and $w_y = (b_y, d_y, u_y)$ be opinions about two distinct propositions x and y , respectively. Then the opinion about $x \wedge y$ is computed using the conjunction operator as $w_{x \wedge y} = w_x \cdot w_y = (b_{x \wedge y}, d_{x \wedge y}, u_{x \wedge y})$ where,

1. $b_{x \wedge y} = b_x b_y$
2. $d_{x \wedge y} = d_x + d_y - d_x d_y$
3. $u_{x \wedge y} = b_x u_y + b_y u_x + u_x u_y$

Deduction operator of SL allows us to associate uncertainty with rules that are used to infer new facts. In our example, $f_1 \text{ AND } f_2 \Rightarrow f_3$ may not be certain. That is, f_3 may not always follow from $f_1 \text{ AND } f_2$. Let shortly refer to $f_1 \text{ AND } f_2$ as θ . The deduction operator uses two conditional opinions about the rule $\theta \Rightarrow f_3$ while deducing opinions for f_3 : i) $w_{f_3|\theta}$ and ii) $w_{f_3|-\theta}$.

The first conditional opinion $w_{f_3|\theta}$ refers to the belief, disbelief, and uncertainty for the proposition that f_3 follows from θ . It is computed using the evidence derived from observations. Given θ holds, if it is observed that f_3 holds as well, this observation is counted as a positive evidence while computing $w_{f_3|\theta}$; but if f_3 does not hold, the observation is taken as a negative evidence. Let r and s refer to the number of positive and negative evidence derived from observations. Then, following (Jøsang and Grandison, 2003), the belief and disbelief in the opinion $w_{f_3|\theta}$ are computed as $r/(r + s + 2)$ and $s/(r + s + 2)$, respectively.

The second opinion $w_{f_3|-\theta}$ refers to the belief, disbelief, and uncertainty for the proposition that f_3 follows from $-\theta$. This opinion is also computed using the evidence derived from observations, as described above. If we set $w_{f_3|\theta}$ as $(0.8, 0, 0.2)$, this means that most of the time f_3 follows from θ , but sometimes this may not hold.

While describing rules, we can use arbitrary propositional logic statements (**AND** , **OR** , and **NOT**). Then, these rules can be used to infer a new fact, given a set of existing facts. For clarity, we have only shown conjunction operator in this paper, other SL operators can be found elsewhere.

4.4.2. Backward Chaining (input opinions are unknown)

We can also infer new opinions about facts using backward chaining, i.e., through abductive reasoning. For instance, let us assume that we have a SL rule indicating that an explosion may lead to high level of dust. When we have a report asserting some opinions about high level of dust, we can use abductive reasoning or backward chaining to infer there might be an explosion resulting in the dust.

More formally, let us assume that f_1 and f_2 refer to facts representing *unusual dust level* and *explosion* at location l and time t , respectively. In our knowledge base, we may have a SL rule relating these facts, i.e., $f_2 \rightarrow f_1$ with conditional opinions $w_{f_1|f_2}$ and $w_{f_1|\neg f_2}$. The conditional opinions can be considered similar to conditional probabilities in a Bayesian network, e.g., $\Pr(\text{dust}|\text{explosion})$ and $\Pr(\text{dust}|\neg\text{explosion})$. Given an opinion w_1 retrieved from a specific source about f_1 , we can infer another opinion w_2 using abduction operator \otimes of SL (Josang, 2008).

4.4.3. Spatio-temporal Relevance and Reasoning with Location

Consider two facts f_1 and f_2 that assert the same statement (e.g., an explosion occurred) but with a small mismatch in its metadata. For instance, the location metadata l_1 associated with fact f_1 is slightly different to l_2 that is associated with fact f_2 . Since the facts essentially agree with each other, the goal of opinion assessment is two-fold: (i) combining opinions in both the facts (because we now have collaborating evidence) and reduce uncertainty, and (ii) assuming that the location l_1 is not same as l_2 , refine the location metadata associated with the facts. For sufficiently small deviations, we may simplify and assume that $l_1 = l_2$ in which case we can apply the SL consensus operator for fusing the opinions tv_1 and tv_2 to compute tv_3 . However, in general, when $l_1 \neq l_2$, traditional SL cannot be used directly. The authors are exploring spatial-aware extensions of SL (coined **saSL**) to address cases like this, with the objective to compute a new location l_3 where the fact (e.g., explosion) could have occurred with high belief. Early results from this work and its application to a localization scenario for cognitive radio networks can be found in (Chakraborty et al., 2012).

Location information in a report can be in the form of latitude and longitude values. In this case, the comparisons of locations in terms of proximity can be done using a set of geographical calculations. On the other hand, it may be more likely to have address information (e.g., street name, area, city etc.) in reports to indicate a specific location. Therefore, in order to identify proximity or spatio-temporal relevance between locations, we may need to do some reasoning. For this purpose, we can use an ontology that determines the relationships between different locations. The linked geographical data² can be used to determine if an area is within or close to another area. Once we determine the relevance between locations appearing in different facts, we can determine if the opinions about these facts are supporting one another or not.

Let us describe this using the following example. We consider that a sniper has killed someone in an area. There are two reports about the sniper's location. The first report is from local

²<http://linkedgeodata.org>

police and indicates with high belief that “the sniper shot from a building x on street y ”. The second report is from a military force conducting operations in the region. The military force also knows that the sniper shot from the building x . However, it does not want to reveal the positions and density of its sensors on the region. Hence, it shares a report after obfuscating (e.g., generalising) the location information. For instance, if the street y is in the area z , it reports that “the sniper shot from a building in area z ”, being sure that this information does not reveal that it has sensors around the building x . When these two reports are received by the system, the locations indicated within the reports are analysed and the relevance between these facts and opinions are determined. That is, the system can figure out that the building x locates in the area z , therefore, the opinions appearing within these reports support each other. Using deduction (forward chaining), we can infer from the first report that the sniper was in the area z , which supports the second report. Similarly, through abduction (backward chaining), the second report supports the first report. That is, if the sniper was at a building on the area z , it is possible that this building is the building x .

5. Fusing of Assessed Opinions

There are different ways of getting various opinions about the same fact. We can get opinions direct from one or more information sources and discount them based on the trustworthiness of their sources or infer opinions through some sort of reasoning mechanism such as deduction (i.e., forward chaining) or abduction (i.e., backward chaining). Therefore, we may expect more than one opinion about the same facts. At this point, we may fuse these opinions using cumulative fusion operator of SL as described in Definition 1 or some other fusion method. On the other hand, some of these opinions may be in conflict. The conflicting opinions may significantly reduce the quality of the fused opinion, especially in the environments with malicious adversaries. In this work, we advocate that conflicts between opinions should be resolved before fusing. In this section, we detail our approach and justify it using a set of simulations.

5.1. Conflicts between Opinions

When we have two or more opinions about the same fact, these opinions may conflict. Let us assume that we have two information sources that provide opinions some how related to the fact f_1 “The road R is bombed”. These sources are Peter and Jane, whose levels of opinion are 0.8 and 0.75, respectively. Peter reports that the fact f_2 “the road is safe” with opinion $(1, 0, 0)$; this opinion is discounted as $(0.8, 0, 0.2)$ based on his trustworthiness. Given that we have a certain rule $hasExplosion(x) \rightarrow \neg safe(x)$, we infer, lets say, the opinion $(0, 0.8, 0.2)$ for f_1 using abduction based on Peter’s report. On the other hand, Jane reports an explosion on the road with opinion $(0.9, 0, 0.1)$; this opinion is discounted as $(0.675, 0, 0.325)$ based on her trustworthiness. Therefore, we have two opinions for f_1 : $(0, 0.8, 0.2)$ and $(0.675, 0, 0.325)$. It is easy to conclude that these opinions are in conflict. The first opinion indicates that most likely there is no explosion on the road while the other implies that there is a significant possibility of explosion.

The conflict indicates that *at least* one of these opinions is misleading. To resolve this conflict, these opinions may be discounted (possibly at different rates), hence the uncertainty within them is increased enough to resolve the conflict. For instance, discounting the first opinion $(0, 0.8, 0.2)$ with 0.0 makes it $(0, 0, 1)$, which absolutely does not conflict with $(0.675, 0, 0.325)$ or any other

opinion, since it does not contain any belief or disbelief, but only uncertainty. Similarly, discounting both of these opinions with 0.5 leads to opinions (0, 0.4, 0.6) and (0.335, 0, 0.665), which are quite uncertain and do not conflict.

In this paper, we argue that conflicts between opinions may serve as evidence for the necessity to increase uncertainty within them. Although there can be more than one way to define conflicts between subjective opinions, in this paper, we introduce Definition 2 where conflicts are defined based on the satisfiability of *beliefs* and *disbeliefs* within opinions. In the rest of this paper, we represent opinions using only belief and disbelief when a coincide notation is required; that is, we use (b, d) instead of (b, d, u) since $u = 1 - b - d$.

Definition 2. Let $O = \{w^0, w^1, \dots, w^n\}$ be opinions about the same proposition, where each opinion $w^i = (b^i, d^i, u^i)$. These opinions are consistent if it is possible to have a valid opinion that can satisfy all of these opinions. An opinion $w^* = (b^*, d^*, u^*)$ can satisfy the opinion $w^i \in O$ iff $b^i \leq b^*$ and $d^i \leq d^*$. Although there can be infinitely many such w^* , the one with the highest uncertainty (i.e., one with the highest u^*) would be

$$(\max(b^0, b^1, \dots, b^n), \max(d^0, d^1, \dots, d^n)).$$

Therefore, there cannot be any opinion that can satisfy all opinions in O if and only if

$$\max(b^0, b^1, \dots, b^n) + \max(d^0, d^1, \dots, d^n) > 1.$$

That is, O is inconsistent iff $\exists w^i, w^j \in O$ s.t. $b^i + d^j > 1$.

The intuition behind the definition of conflicts between opinions is as follows. Let the ground truth about a proposition x be represented as an opinion $w_x^r = (b_x^r, d_x^r)$. Also, let $w_x^i = (b_x^i, d_x^i, u_x^i)$ be an arbitrary opinion about x . If w_x^i has a higher *belief* than w_x^r does (i.e., $b_x^r < b_x^i$), then w_x^i is misleading. Similarly, if w_x^i has a higher *disbelief* than w_x^r does (i.e., $d_x^r < d_x^i$), then w_x^i is misleading. How much w_x^i is misleading depends on how much extra belief and disbelief it imposes. On the other hand, if $b_x^i < b_x^r$ and $d_x^i < d_x^r$, then w_x^i is not misleading, because it does not impose any extra belief or disbelief, but only extra uncertainty that conflicts with neither the *belief* nor the *disbelief* within the ground truth w_x^r . In reality, we do not have the ground truth about x . Therefore, we cannot say whether w_x^i is misleading or not. However, if we have another opinion $w_x^j = (b_x^j, d_x^j, u_x^j)$, we can reason about if it is possible to have a ground truth for which neither w_x^i nor w_x^j is misleading. Such a ground truth exists if and only if $\max(b_x^i, b_x^j) + \max(d_x^i, d_x^j) \leq 1$. If such a ground truth cannot exist, we say w_x^i and w_x^j are in conflict, i.e., at least one of them must be misleading at some degree.

5.2. Resolving Conflicts

Before fusing opinions regarding a specific fact, for example using cumulative fusion operator in Definition 1, the conflicts between opinions should be resolved. Once we determine conflicting opinions, several approaches can be used for this purpose. In this section, we introduce three conflict resolution approaches and demonstrate their performance in terms success of the fusion operation. In these approaches, discounting operator is used to resolve conflicts between opinions by increasing uncertainty in some of these opinions. Let us have two conflicting opinions w_i and w_j from information sources i and j with opinion values t_i and t_j , respectively. These opinions are discounted by coefficients $0 \leq x_i, x_j \leq 1$ based on information like t_i and t_j .

1. **Trust-based deleting:** If two opinions w_i and w_j are in conflict, the opinion from the less trustworthy source is deleted, and if both sources are equally trustworthy both opinions are deleted. Thus, if the trust we have in the source of opinion w_i is greater than that of the source of w_j ($t_i > t_j$) then $x_i = 1$ and $x_j = 0$, and when $t_1 = t_2$, we assign $x_1 = x_2 = 0$.
2. **Trust-based discounting:** If two opinions w_i and w_j are in conflict, they are discounted in proportion to the trustworthiness of their sources. That is, the more trust worthy opinion is discounted relatively less. Let us define a conflict over $w_i = (b_i, d_i, u_i)$ and $w_j = (b_j, d_j, u_j)$ as $b_i + d_j > 1$. Then, the additional discounting factors for w_i and w_j are computed as $t_i/(b_i t_i + d_j t_j)$ and $t_j/(b_i t_i + d_j t_j)$, respectively.
3. **Evidence-based discounting:** The previous two conflict resolution methods uses discounting based on trust to resolve conflicts. Trustworthiness of information sources are computed based on past experience, i.e., evidence. Evidence relates to how reliable or unreliable a source was in the past. The evidence-based discounting uses TRIBE (Sensoy et al., 2013) approach for conflict resolution and depends on evidence analysis instead of trust values. As described before, each opinion in the proposed system is already discounted by the trustworthiness of its source. Discounting an opinion further to resolve a conflict implies making its trustworthiness less than the trustworthiness of its source. For instance, let us assume that our trust for Jack and Jane as information sources are 0.8 and 0.9, respectively; and they give opinions (1, 0, 0) and (0, 1, 0) for the fact “the road is not safe”, respectively. After discounting based on the trustworthiness of their sources we have opinions: (0.8, 0, 0.2) and (0, 0.9, 0.1). These opinions are still in conflict. To resolve the conflict, we need to discount them further. For instance, we can discount both of these opinions with 0.5 and have two non-conflicting opinions: (0.4, 0, 0.6) and (0, 0.45, 0.55). Hence, the trustworthiness of the first opinion becomes $0.5 \times 0.8 = 0.4$ and that of the second becomes 0.45. Let us assume that the trust in the source of the first opinion is computed as $(r+1)/(r+s+2) = 0.8$ based on Equation 1 where $r = 8$, $s = 0$. We can trivially calculate that 11 extra negative evidence, i.e., speculative evidence, should be added to s to decrease the trust from 0.8 to 0.4, i.e., $(8+1)/(8+11+2) = 0.4$. TRIBE aims at optimising the total amount of speculative evidence necessary to resolve all conflicts by finding the best discounting factors. Therefore, it considers not only the trust values, but also the evidence used to calculate them.

We evaluated the performance of the listed conflict resolution approaches in cumulative fusion through a set of simulations. In these simulations, there are two types of information sources that provides opinions about fact: honest and malicious sources. At each time step, a fact is randomly selected and a ground truth is generated for the fact. Our fusion model collects opinions provided by various information sources and fuse these opinions by i) detecting conflicts, ii) resolving conflicts, and iii) applying the cumulative fusion operator to generate a single fused opinion. The success of the fusion operation is measured using absolute error. Let (b, d, u) be the ground truth and (b', d', u') be the fused opinion, then the *absolute error* is computed as $err_{B(a)} = abs(\delta_b) + abs(\delta_d)$, where $\delta_b = b' - b$ and $\delta_d = d' - d$.

The honest sources provide opinions close to ground truth; they add small amount of Gaussian noise with $N(0, 0.01)$ to belief and disbelief within the ground truth. The malicious sources behave like honest sources for a while until they build up some trust (above 0.7) and then they start to provide misleading opinions. The misleading opinions are generating by switching the belief and disbelieve within genuine opinions, which are close to the ground truth. In order to increase the uncertainty in the behaviour of sources, we also let honest sources provide misleading opinions with a probability of 0.1. During simulations there are 10 information sources

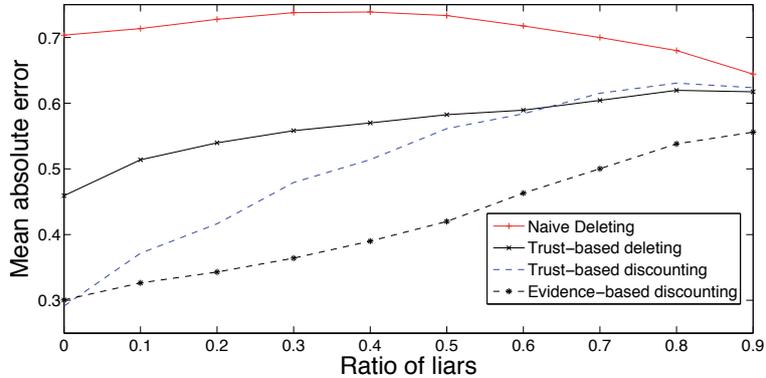


Figure 5: Error of cumulative fusion operator in the face of liars for varying R_{liar}

and we have run the simulations for different ratio of liars. For each ratio of liars, we repeat our experiments 10 times and we demonstrate the mean absolute error for our experiments in Figure 5. In the figure, *Naive Deleting* method is a baseline method in our evaluations. This method simply deletes all of the conflicting opinions before fusion. Our results indicate that trust-based discounting and evidence-based discounting have similar performance for low ratio of liars. However, considering only trust is not enough for higher ratios of liars. Trust-based deleting and naive deleting approaches do not perform good. They lead to highly erroneous fused opinions. Evidence-based discounting is by far the best conflict resolution approach. It allows misleading opinions to be discounted before the cumulative fusion takes place. Therefore, we suggest to use evidence-based discounting in the proposed system.

6. Experimental Evaluation

The SL-based opinion assessment process has been implemented on the *Information Fabric* (or *Fabric*) that serves as the underlying platform for executing the opinion assessment workflows. The opinion assessment operations appear as a service that is accessible over the Fabric. Specifically, packaged as a Java library, the SL-based operators are invoked as needed by the opinion assessment service.

The SL-based opinion assessment has also been implemented using Controlled English as the inference engine. opinion assessment framework appears as a plugin in the Controlled English Store. On execution of the inference engine, the opinion assessment plugin will be invoked and will perform the assessment. Subjective logic operator has been used to represent the rules with opinions hence associating rules with some uncertainty. This complies to real life scenario and hence a good experimentation of the assessment.

All our experiments were performed on the SYNCOIN dataset (Graham et al., 2011). We used the CE Store to a priori extract facts from the dataset. These facts are then streamed over the Fabric to an in-memory database wherein facts are annotated with opinion metadata.

6.1. Performance

Table 6 summarizes the performance overhead added by our opinion assessment approach. We compared the latency and throughput of a nop (no operation) Fabric service with that of

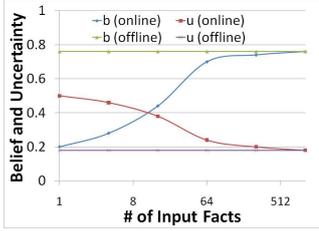


Figure 6: Forward Chaining

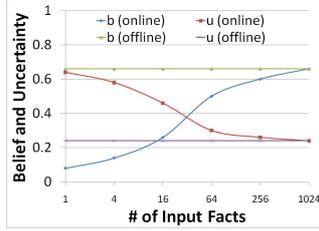


Figure 7: Consensus (Self Chaining)

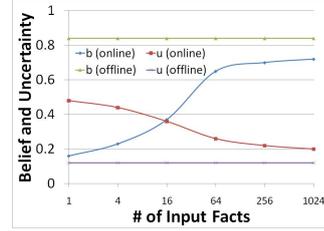


Figure 8: Backward Chaining

Metric	Mean	Std dev
Throughput (per sec)	33.6 / 33.5 / 2.98%	2.4 / 2.8 / 16.67%
Latency (ms)	32.2 / 36.4 / 13.04%	12.1 / 13.2 / 9.09%

Figure 9: Performance Overhead $x/y/z$ where x denotes the performance of a `nop` Fabric service, y denotes the performance of our opinion service and z shows the relative overhead (i.e., percentage decrease in throughput and percentage increase in latency)

our Fabric opinion service by streaming events through them at high speed. The `nop` Fabric service takes a stream of events as input and returns an identical stream of events as output. Thus, execution of `nop` reflects the cost of an event entering and exiting the Information Fabric platform without any processing within the Fabric platform. The trust service takes a stream of events as input and returns a opinion annotated stream of events as output. We have repeated the experiment by calling each of the SL operators shown in Table 1. The variation in latency and throughput across these SL operators were hardly noticeable. All our experiments were repeated seven times. The experiments were repeated until we attained reasonable variance and 5-7 runs reduced the variance.

Our initial results indicate that including a trust service does not significantly alter the throughput of the Fabric platform. In general, we expect the information analytics operator is likely to be at least as complex (if not more) than the opinion assessment operator; hence, the relative loss in throughput is expected to be reasonably small (roughly 3%). With regard to latency, the trust service adds a relatively larger overhead (13%). This is because in our current implementation, the opinion assessment operator does not run concurrently with the analytics operator; we require the analytics operation to be complete and the output information product be available before the opinion assessment operation can infer the opinion metadata on the information product. For future implementations, we anticipate that once the choice of the analytics operation is determined, e.g., which CE rule will be applied, we could concurrently issue a request to the opinion service without waiting for the analytics operation to complete. Hence, in this case, we expect that the latency overhead due to the trust service would decrease.

6.2. Accuracy

Next, we compare the accuracy of the assessed opinion using our online opinion assessment system versus an offline approach. The offline approach waits for all reports to be received before assessing opinion in various facts. In contrast, our online approach operates on a stream of incoming reports with the goal of ascertaining opinion in various facts as information arrives. Assuming that opinion levels in reports do not change with time, the opinion assessed by the

offline approach can serve as a benchmark, representing a ceiling for the opinion levels that can be ascertained as the entire body of information is available at decision time. Against this benchmark, we can assess how our online approach fares. The online approach operates on a compact set of facts where the incoming reports are discarded immediately after fact extraction and opinion assessment. We will refer to the difference between the opinion levels attained by the offline and online approaches as the *accuracy* in the opinion assessment.

Figures 6, 7 and 8 compare the accuracy for different classes of opinion analytics operators. Note that the forward chaining operators (such as the SL AND and OR operators) and the SL consensus operator are associative and commutative and, hence, the order in which these operators are applied does not change the outcome. As a result, the online opinion assessment approach will eventually converge to the value attained by the offline approach when these operators are used; this behavior is seen in Figures 6 and 7. On the other hand, backward chaining operators (such as the SL deduction and abduction operators) are neither commutative nor associative and, hence, the order of their application does matter. In the backward chaining case, see figure 8, the online opinion assessment system exhibits inaccuracies by assessing opinion at levels that remain “distant” to those attained by the offline approach.

The inaccuracy notwithstanding, the online approach allows for fast-paced evaluation of opinion as information arrives. In the particular example, without the computer processing time, there is practically zero delay from the moment the information arrives until a new opinion level can be assessed. When opinion level reach satisfactory levels and remain at such levels for a sufficient period of time, decisions could be made and actions taken at earlier times, resulting in fast-paced execution of the pertinent OODA loop as well. Clearly, though, the level in opinion accuracy attained and the responsiveness with which it is attained presents a design and system management trade-off parameter. To this end, in future work, we will explore operational implications of alternative designs where, for example, we trade-off delay to accuracy by operating in a near-streaming manner. In the latter case, the system could operate over time frames assessing opinion based on reports arriving over a single frame in an offline fashion, but assessing opinion over successive windows in an online fashion.

7. Related Work

There has been research in areas of trust based data management. Patwardhan *et al.* propose a trust-based data management framework for enabling individual devices to harness the potential power of distributed computation, storage, and sensory resources available in pervasive computing environments (Patwardhan et al., 2005). They take a holistic approach that considers trust, security, and privacy issues of data management in these environments.

Lim *et al.* has explored data stream management and exploited the notion of confidence policy while taking into account trustworthiness of data items in data management and query processing. They propose a provenance-based framework that enforces confidence policies in the evolution of continuous queries over streaming data. Based on the notion of physical and logical networks, they introduce the notions of the physical and logical provenances for data items, respectively. The authors also introduce a cyclic framework of computing actual trust scores of data items and network nodes based on the value and provenance similarities embedded in data items.

Unlike these approaches, in our work, we use unstructured reports from unreliable information sources and derived structured knowledge based on Controlled English using ITA (International Technology Alliance) assets. ITA (International Technology Alliance) is a collaborative research

alliance between UK Ministry of Defence (UK MoD) and US Army Research Laboratory (ARL) and a consortium of leading industry and academic partners. Some of the assets that has resulted due to this research collaboration includes Controlled English (Mott, 2010) and Information Fabric (Wright et al., 2009). Controlled English (CE) is an ITA asset that has been extended and modified by the ITA programme. Mott defines the syntax with examples of CE formatted text and shows how to map to predicate logic (Mott, 2010). CE is a type of Controlled Natural Language (CNL) and is mainly used to be readable by any one who knows English. It represents information in a very structured form using certain syntax. This syntax can be parsed and the content can be understood by a computer. Hence, CE enables the communication between computer and humans.

Schwitter has described the use of controlled natural language as a high level specification language for modelling commonsense reasoning problems (Schwitter, 2015). The paper shows how defaults and exceptions can be incorporated into a current controlled natural language and what is required to signify and motive with them in a not so straight way.

In this work, we use a trust model that is based on trust evidence. The retrieval or derivation of trust evidence is out of the scope of this paper. Usually, there are two type of evidence: direct and indirect. In the literature, there are several trust approaches where direct evidence is combined with indirect evidence to model trust in information sources. Direct evidence is based on personal observations, while indirect evidence is received from others. Jøsang and Ismail proposed the beta reputation system (BRS) (Jøsang and Ismail, 2002). It estimates the trustworthiness of an information source using beta probability density functions, as we do in this paper. For this purpose, aggregation of direct evidence and indirect evidence from information sources are used as the parameters of beta distributions. Evidence shared by sources are equivalent to binary opinions in Subjective Logic (Jøsang, 2011). Whitby *et al.* extended BRS to handle misleading indirect evidence from malicious agents using a majority-based algorithm (Whitby et al., 2005). Teacy *et al.* proposed TRAVOS (Teacy et al., 2006), which is similar to BRS, but it uses personal observations about information sources to filter misleading indirect evidence.

Subjective logic has been extended to include belief updates from partially visible evidence (Kaplan et al., 2015).. Lance et.al show that assets of the partial observable update as a function of the state likelihood and demonstrate the use of these likelihoods for a trust estimation application. The value of the partial observable updates is shown through different imitations including the trust estimation case.

Yu and Singh proposed a trust approach that handles misleading indirect evidence using a version of weighted majority algorithm (Yu and Singh, 2003). In their algorithm, weights are assigned to information sources. These weights are initiated as 1.0 and can be considered as the trustworthiness of the corresponding sources. The algorithm makes predictions about trust related propositions (e.g, i is trustworthy) based on the weighted sum of indirect evidence (i.e., ratings) provided by those sources. The authors proposed to tune the weights after an unsuccessful prediction so that the weights assigned to the unreliable sources are decreased. They assume that the ratings from dishonest sources may conflict with the personal observations. By decreasing the weights of these sources over time, misleading evidence is filtered.

In this paper, we describe conflicts between binomial opinions and analyse evidence about information sources to resolve conflicts between opinions before performing fusion. Conflicts in knowledge lead to inconsistencies that hamper the reasoning over the knowledge. Therefore, before using such knowledge bases, their conflicts should be resolved. Gobeck and Halaschek (Golbeck and Halaschek-Wiener, 2009) present a belief revision algorithm for ontologies, which is based on trust degrees of information sources to remove conflicting statements from a knowl-

edge base. However, as the authors point out, the proposed algorithm is not guaranteed to be optimal. Dong *et al.* (Dong et al., 2009) propose to resolve conflicts in information from multiple sources by a voting mechanism. Double counting in votes is avoided by considering the information dependencies among sources. The dependences are derived from Bayesian analysis.

When considering multiple sources espousing multiple claims, it is possible to estimate their reliability through *corroboration* without direct and/or indirect evidence. For example, fact-finding algorithms aim to identify the *truth* given conflicting claims. Yin *et al.* proposed TruthFinder (Yin et al., 2007) which utilizes an iterative approach to estimate trustworthiness of information sources and information they provide. Their approach based on the assumption that a source is trustworthy if it provides many pieces of true information, and a piece of information is likely to be true if it is provided by many trustworthy sources. Therefore, very similar to BRS, TrustFinder also assumes that the information provided by the majority is trustworthy.

In our paper, the key differences compared to the papers described in this section are mainly the streaming nature of updates (online mode). Most facts are inserted as information, as inferences are derived we can see that we are propagating from source to information to inference and a trust annotation for the inference. Unlike our online mode, information is non-streaming in the offline mode.

8. Conclusions

Effective applications of the OODA loop requires a fast-paced execution of its *Observe* and *Orient* phases. The latter becomes challenging when faced with voluminous information of questionable provenance and trustworthiness arriving in streaming fashion from multiple source, as the case might be in military coalition environments.

In this paper, we introduced an opinion assessment framework comprising a family of opinion operators, based on subjective logic. The framework is suitable for dynamic information flows assessing opinions for data-in-motion online rather than for a large corpus of static data. We have also described an implementation of our opinion assessment framework using ITA assets (Information Fabric Services (Wright et al., 2009) and Controlled English Fact Store (Mott, 2010)) and presented an experimental evaluation that quantifies the efficacy (e.g., accuracy and overhead) of our framework. Initial analysis shows promising results compared to a benchmark offline design when utilizing opinion assessment operators that satisfy associative and commutative properties. If these properties do not hold, the online assessment exhibit inaccuracies and considerations are under way to improve accuracy in this case as well.

In future work, we also plan to explore additional opinion assessment frameworks that may be even better suited at handling dynamic information flows, for example, considering performing opinion assessment across a loosely coupled offline and online system.

Acknowledgements

Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the 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 authorised to reproduce and distribute reprints for Government purposes

notwithstanding any copyright notation hereon. Dr. Şensoy thanks the U.S. Army Research Laboratory for its support under grant W911NF-13-1-0243 and The Scientific and Technological Research Council of Turkey (TUBITAK) for its support under grant 113E238.

References

- Bisdikian, C., Kaplan, L. M., and Srivastava, M. B. (2011). Quality of information in sensor networks. In *IBM Res. Rep. RC25254*.
- Bisdikian, C., Kaplan, L. M., Srivastava, M. B., Thornley, D. J., Verma, D., and Young, R. I. (2009). Building principles for a quality of information specification for sensor information. In *12th Int'l Conf. on Information Fusion (FUSION 2009)*.
- Brehmer, B. (2005). The dynamic ooda loop: Amalgamating boyd's ooda loop and the cybernetic approach to command and control. In *Proceedings of the 10th international command and control research technology symposium*.
- Chakraborty, S., Srivastava, M. B., Bisdikian, C., and Ganti, R. (2012). Localization in cognitive radio systems in the presence of spatially obfuscated data. Technical Report TR-UCLA-NESL-201205-01, NESL.
- Dong, X. L., Berti-Equille, L., and Srivastava, D. (2009). Integrating conflicting data: The role of source dependence. In *Proc. of the 35th International Conference on Very Large Databases*, Lyon, France.
- Golbeck, J. and Halaschek-Wiener, C. (2009). Trust-based revision for expressive web syndication. *Journal of Logic and Computation*, 19(5):771–790.
- Graham, J. L., Rimland, J., and Hall, D. L. (2011). A COIN-inspired synthetic dataset for qualitative evaluation of hard and soft fusion systems. In *14th Int'l Conf. on Information Fusion (FUSION 2011)*, Chicago, IL, USA.
- Jøsang, A. (2001). A logic for uncertain probabilities. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 9(3):279–311.
- Jøsang, A. (2002). The consensus operator for combining beliefs. *Artificial Intelligence*, 141(1/2):157–170.
- Jøsang, A. (2007). Probabilistic logic under uncertainty. In *Proceedings of the thirteenth Australasian symposium on Theory of computing*, pages 101–110.
- Josang, A. (2008). Conditional reasoning with subjective logic. *Journal of Multiple-Valued Logic and Soft Computing*, 15(1):5–38.
- Jøsang, A. (2011). *Subjective Logic*. Book Draft.
- Jøsang, A. and Grandison, T. (2003). Conditional inference in subjective logic. In *Proceedings of the 6th International Conference on Information Fusion*.
- Jøsang, A. and Ismail, R. (2002). The beta reputation system. In *Proc. of the 15th Bled Electronic Commerce Conference e-Reality: Constructing the e-Economy*, pages 48–64.

- Jøsang, A. and McAnally, D. (2005). Multiplication and comultiplication of beliefs. *Int. J. Approx. Reasoning*, 38(1):19–51.
- Jøsang, A. and Presti, S. L. (2004). Analysing the relationship between risk and trust. In *Proceedings of the Second International Conference on Trust Management (iTrust)*, Oxford.
- Kamvar, S., Schlosser, M., and Garcia-Molina, H. (2003). EigenTrust: Reputation management in P2P networks. In *WWW Conference*.
- Kaplan, L., Şensoy, M., Chakraborty, S., and de Mel, G. (2015). Partial observable update for subjective logic and its application for trust estimation. *Information Fusion, Elsevier (online)*.
- Kuter, U. and Golbeck, J. (2007). SUNNY: A New Algorithm for Trust Inference in Social Networks, using Probabilistic Confidence Models. In *AAAI-07*.
- Mott, D. (2010). Summary of ita controlled english. ITA Tech. Report. retrieved from “<https://www.usukita.org/papers/5658/details.html>,” May 4, 2012.
- Patwardhan, A., Perich, F., A. Joshi, T. F., and Yesha, Y. (2005). Active collaborations for trustworthy data management in ad hoc networks. In *Proceedings of the 2nd IEEE International Conference on Mobile Ad-Hoc and Sensor Systems*.
- Schafer, J. B., Konstan, J., and Riedl, J. (1999). Recommender Systems in E-Commerce. In *ACM Conference on Electronic Commerce*.
- Schwiter, R. (2015). Handling defaults and their exceptions in controlled natural language. *Language Production, Cognition, and the Lexicon, Springer*, 48(4):313–330.
- Sensoy, M., Bisdikian, C., Oren, N., Burnett, C., Norman, T. J., Srivastava, M. B., and Kaplan, L. (2012). Trust and obfuscation. In *SPIE Defense, Security, and Sensing Symposium*, Baltimore, MD, USA.
- Sensoy, M., de Mel, G., Pham, T., Kaplan, L., and Norman, T. J. (2013). Tribe: Trust revision for information based on evidence. In *Proceedings of 16th International Conference on Information Fusion*, pages 914–921.
- Teacy, W., Patel, J., Jennings, N., and Luck, M. (2006). TRAVOS: Trust and reputation in the context of inaccurate information sources. *Autonomous Agents and Multi-Agent Systems*, 12(2):183–198.
- Wang, S., Su, L., Li, S., Yao, S., Hu, S., Kaplan, L., Amin, M. T., Abdelzaher, T., and Wang, H. (2015). Scalable social sensing of interdependent phenomena. In *14th ACM/IEEE Conference on Information Processing in Sensor Networks (IPSN)*, Seattle, WA.
- Whitby, A., Jøsang, A., and Indulska, J. (2005). Filtering out unfair ratings in Bayesian reputation systems. *The Icfaïn Journal of Management Research*, 4(2):48–64.
- Wright, J., Gibson, C., Bergamaschi, F., Marcus, K., Pressley, R., Verma, G., and Whipps, G. (2009). A dynamic infrastructure for interconnecting disparate isr/istar assets (the ita sensor fabric). In *12th Int’l Conf. on Information Fusion (FUSION 2009)*, Seattle, WA, USA.

- Xiong, L. and Liu, L. (2004). Supporting reputation based trust in peer-to-peer communities. In *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, vol. 71, 16(7).
- Yin, X., Han, J., and Yu, P. S. (2007). Truth discovery with multiple conflicting information providers on the web. In *Proceedings of the Conference on Knowledge and Data Discovery*.
- Yu, B. and Singh, M. (2003). Detecting deception in reputation management. In *Proceedings of Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 73–80.