Obfuscation of Semantic Data: Restricting the Spread of Sensitive Information

Federico Cerutti[1,⋆], Geeth R. de Mel[2], Timothy J. Norman[1], Nir Oren[1], Artemis Parvizi[1], Paul Sullivan[3], and Alice Toniolo[1]

[1] Department of Computing Science, University of Aberdeen, UK,
{f.cerutti,t.j.norman,n.oren,a.parvizi,a.toniolo}@abdn.ac.uk
[2] IBM TJ Watson Research Center, NY, USA, grdemel@us.ibm.com
[3] INTELPOINT Incorporated, PA, USA, paul.sullivan@intpt.net

Abstract. This paper investigates how to restrict the spread of sensitive information. This work is situated in a military context, and provides a tractable method to decide what semantic information to share, with whom, and how. The latter decision is supported by obfuscation, where a related concept or fact may be shared instead of the original. We consider uncertain information, and utilize Subjective Logic as an underlying formalism to further obfuscate concepts, and reason about obfuscated concepts.

Keywords: semantic knowledge, strategic interaction, knowledge obfuscation

1 Introduction

Controlling the spread of sensitive information is a problem in various contexts. In everyday life, users share information across a plethora of social networks and other media. This raises concerns about unwanted usage of such information [1, 2]. In strategic contexts [3, 4, 5], information sharing can have serious economic or life threatening repercussions. This paper therefore examines how information sharing can take place while minimizing its negative impact through the use of *obfuscation*.

Following [6], a technique for restricting the spread of information to the "desired" audience only is referred to as *obfuscation*, defined as:

> Information-obfuscation (or data-masking) is the practice of concealing, restricting, fabricating, encrypting, or otherwise obscuring sensitive data. [6]

While other approaches [1, 2, 6] focused on obfuscation of *quantitative* information — e.g. accelerometer data from a smartphone — we focus on *qualitative* information linked through a semantic description of the domain. The main contribution of this paper is an innovative and sound ontology based obfuscation technique, useable in non-cooperative environments.

We consider an implicit exchange of information between two agents: a $sender$ and a $receiver$. The $sender$ wants the $receiver$ to know some pieces of information, and at the same time it wants to keep some inferences[4] that could be surmised from this

[⋆] Corresponding author.

[4] In this paper, *inference* refers to any reasoning process that — when applying a specific rule or rules (e.g. *deductive modus ponens*) — leads to a conclusion given a set of premises.

information private. This is analogous to the work described in [2] where users wish to share the *activity level* obtained from their smartphone accelerometer with an app for medical advice, while keeping the specific *activity type* private.

In this paper we present the Semantic Obfuscation Framework (SOF), which adopts the $sender$'s point of view, and thus starts considering (1) its domain model (ontology), (2) the information to be shared, and (3) the information to be kept private. Then it identifies the relevant ontological relationships between the shared and the private information, and computes, using Subjective Logic[5] (SL) [7], the likelihood that the $receiver$ knows it thus leading to deriving the private information. Note that the use of SL in such a context is not novel: SL is utilized in [8] for evaluating source of ontological information, and in [9, 10] for computing trust/reputation degrees with uncertainty.

The paper is organized as follows. Section 2 shows a realistic military scenario and the the main concepts of SL. Section 3 presents the requirements necessary for applying our proposal for obfuscating semantic data, which is then illustrated in Section 4. Finally, Section 5 concludes the paper by discussing related and future work. Proofs are omitted or sketched due to space constraints.

2 Motivation and Background

A realistic military scenario [11] is used throughout this paper as a running example for demonstrating our proposal. Such a scenario (§ 2.1) is formalized using the Intelligence, Surveillance, Target Acquisition and Reconnaissance (ISTAR) ontology O_I [12, 13][6].

2.1 A Motivating Scenario

In [11], a realistic military scenario is developed involving peace-keeping in the country of Sincor where coalition forces successfully executed a campaign to liberate it from a dictatorial regime. A military task force is made up of members from many different nations working together as a coalition, including local forces. The coalition has divided responsibility for the country into different sectors, with a field grade officer in charge of each sector. The Combined Joint Task Force (CJTF) commander is in critical need of a capability to know the movement of all Person of Interest (POI) and their activities within the country and to have full situational awareness of the activities of the members of various Violent Extremist Organizations (VEO) that work to destabilize the country and push it back to war.

Last night, two American diplomats were kidnapped, and their current whereabouts are unknown. The main VEO suspected of the kidnapping is Sumer, which operates from a safe haven in the hills just outside of Kish. The CJTF Commander "stands-up" a Crisis Action Team (CAT) to help manage the fluid situation. Intelligence has contacted other intelligence organizations within the coalition to try and determine the exact whereabouts of the hostages and the size of the force holding them. CJTF Intelligence is also preparing intelligence products for the CAT on the kidnapping and any

[5] Subjective Logic extends probability theory by expressing uncertainty about the probability values themselves.

[6] `http://goo.gl/feTio9`

coalition forces in the area as well as the available capabilities of coalition local partners. This information will be used by members of the CAT in formulating possible Courses of Actions (COA) for the commander to consider.

A few hours ago, Intelligence determined a POI who is very likely to be involved in the kidnapping. The CJTF Commander needs to contact the coalition local partners because he will need support for constant surveillance of such a POI, but only if it is unlikely that they will infer that this is a hostage rescue operation: the coalition local partners might otherwise jeopardize the operation due to insufficient training.

Therefore, the CJTF Commander asks Intelligence to evaluate the likelihood of inferences that could be made by coalition local partners. In particular, he informs Intelligence that he needs to keep the nature of the operation (i.e. hostage rescue) confidential. Intelligence, following the approach presented in this paper, replies that the probability that the coalition partners will infer the nature of the operation is 12%, a value that the CJTF considers acceptable. Thus the CJTF asks the coalition local partners to support its operation with constant surveillance of a POI.

2.2 Overview of Subjective Logic

To assess the uncertainty in reasoning about another agent's knowledge and, ultimately, to derive a metric of obfuscation (§ 4), we rely on Subjective Logic (SL) [7], which extends probability theory by expressing uncertainty about the probability values themselves. Like Dempster-Shafer evidence theory [14, 15], SL operates on a *frame of discernment* — i.e. a set of *atomic or primitive states*, namely variable assignments, e.g. if d represents the possible results of rolling a dice, atomic states are $d = 1$, $d = 2$, …, $d = 6$ — denoted by Θ.

Definition 1 (Belief Mass Assignment). *Given a frame of discernment Θ, we can associate a* belief mass assignment (BMA) $m_\Theta(x)$ *with each substate $x \in 2^\Theta$ such that*

1. $m_\Theta(x) \geq 0$
2. $m_\Theta(\emptyset) = 0$
3. $\sum_{x \in 2^\Theta} m_\Theta(x) = 1$

When we speak of belief in a certain state, we refer not only to the belief mass in the state, but also to the belief masses of the state's substates. Similarly, when we speak about disbelief, that is, the total belief that a state is not true, we need to take substates into account. Finally, SL also introduces the concept of uncertainty, that is, the amount of belief that might be in a superstate or a partially overlapping state. These concepts can be formalized as follows.

Definition 2 (Belief/Disbelief/Uncertainty Functions/Relative atomicity/Probability Expectation). *Given a frame of discernment Θ, a belief mass assignment m_Θ on Θ, and a state x, we define:*

- *the* belief function *for x:* $b(x) = \sum_{y \subseteq x} m_\Theta(y), \ \ x, y \in 2^\Theta$*;*
- *the* disbelief function *for x:* $d(x) = \sum_{y \cap x = \emptyset} m_\Theta(y), \ \ x, y \in 2^\Theta$*;*

- *the* uncertainty function *for x:* $u(x) = \sum_{\substack{y \cap x \neq \emptyset \\ y \not\subseteq x}} m_\Theta(y), \;\; x, y \in 2^\Theta$*;*
- *the* relative atomicity *for x w.r.t. $y \in 2^\Theta$:* $a(x/y) = \dfrac{|x \cap y|}{|y|}, \;\; x, y \in 2^\Theta, y \neq \emptyset$*;*
- *the* probability expectation *for x:* $\mathbb{E}[x] = \sum_y m_\Theta(y) a(x/y), \;\; x, y \in 2^\Theta$.

In particular, let us consider a *focused frame of discernment*, viz. a binary frame of discernment containing a state and its complement.

Definition 3 (Focused Frame of Discernment/Focused BMA/Focused Relative Atomicity). *Given $x \in 2^\Theta$, the frame of discernment denoted by $\tilde{\Theta}^x$, which contains two atomic states, x and $\neg x$, where $\neg x$ is the complement of x in Θ, is the* focused frame of discernment *with focus on x. Let $\tilde{\Theta}^x$ be the focused frame of discernment with focus on x of Θ. Given a belief mass assignment m_Θ and the belief, disbelief and uncertainty functions for x ($b(x)$, $d(x)$ and $u(x)$ respectively), the* focused belief mass assignment*, $m_{\tilde{\Theta}^x}$ on $\tilde{\Theta}^x$ is defined as*

$$\begin{aligned} m_{\tilde{\Theta}^x}(x) &= b(x) \\ m_{\tilde{\Theta}^x}(\neg x) &= d(x) \\ m_{\tilde{\Theta}^x}(\tilde{\Theta}^x) &= u(x) \end{aligned}$$

The focused relative atomicity *of x (which approximates the role of a prior probability distribution within probability theory, weighting the likelihood of some outcomes over others) is defined as*

$$a_{\tilde{\Theta}^x}(x/\Theta) = [\mathbb{E}[x] - b(x)]/u(x)$$

For convenience, and when clear from the context, the focused relative atomicity of x is abbreviated to $a(x)$.

An opinion consists of the belief, disbelief, uncertainty and relative atomicity as computed over a focused frame of discernment.

Definition 4 (Opinion). *Given a focused frame of discernment Θ containing x and its complement $\neg x$, and assuming a belief mass assignment m_Θ with belief, disbelief, uncertainty and relative atomicity functions on x in Θ of $b(x)$, $d(x)$, $u(x)$, and $a(x)$, we define an* opinion *over x, written o_x as*

$$o_x \equiv \langle b(x), d(x), u(x), a(x) \rangle$$

We denote the set of all possible SL opinion 4-ples with $\mathbb{O} \subseteq [0, 1]^4$.
The probability expectation of an opinion is denoted as $\mathbb{E}[o_x] = b(x) + u(x) \cdot a(x)$.

Given two opinions about propositions x_1 and x_2, [7] defines a conjunction operator as follows.

Definition 5 (Propositional Conjunction). *Let* $o_{x_1} = \langle b(x_1), d(x_1), u(x_1), a(x_1) \rangle$ *and* $o_{x_2} = \langle b(x_2), d(x_2), u(x_2), a(x_2) \rangle$ *be opinions about* x_1 *and* x_2*. Let* $o_{x_1 \wedge x_2} = \langle b(x_1 \wedge x_2), d(x_1 \wedge x_2), u(x_1 \wedge x_2), a(x_1 \wedge x_2) \rangle$ *be the opinion such that:*

$$
\begin{aligned}
b(x_1 \wedge x_2) &= b(x_1)\, b(x_2) \\
d(x_1 \wedge x_2) &= d(x_1) + d(x_2) - d(x_1)\, d(x_2) \\
u(x_1 \wedge x_2) &= b(x_1)\, u(x_2) + u(x_1)\, b(x_2) + u(x_1)\, u(x_2) \\
a(x_1 \wedge x_2) &= \frac{b(x_1)\, u(x_2)\, a(x_2) + u(x_1)\, a(x_1)\, b(x_2) + u(x_1)\, a(x_1)\, u(x_2)\, a(x_2)}{b(x_1)\, u(x_2) + u(x_1)\, b(x_2) + u(x_1)\, u(x_2)}
\end{aligned}
$$

3 Requirements for Semantic Obfuscation

We now turn our attention to the requirements for providing a formal definition for an obfuscation procedure. From a general perspective, we consider two agents: the $sender$ shares *White knowledge* — i.e. a piece of information somehow linked to a domain ontology — with a $receiver$. The $sender$ wants to keep private some pieces of information — *Black knowledge* — that might be surmised from the White knowledge by exploiting semantic connections. These connections are derived from the $sender$'s domain ontology which we assume has been built by an ontology engineer.

From a formal perspective, if not explicitly mentioned, we refer to an arbitrary but fixed ontology O built in the monotonic description logic language $\mathcal{EL}$ [16]. O is a finite set of axioms in $\mathcal{EL}$; N_C^O is the set of *concept names*; N_R^O is the set of *role names*; N_I^O is the set of *individual names*. Concepts $v_2, v_3 \in N_C^O$ can be inductively composed with the following constructs: $\top \mid \bot \mid v_1 \mid \{a\} \mid v_2 \sqcap v_3 \mid \exists r.v_2$ where $\top$ denotes the top concept; $\bot$ denotes the bottom concept; $v_1 \in N_C^O$; $a \in N_I^O$; $r \in N_R^O$. $\mathcal{EL}$ supports concept inclusion $v_2 \sqsubseteq v_3$ and membership $v_1(a)$. An axiom α is an assertion in $\mathcal{EL}$ which is a well-formed formula; $TBox$ represents all axioms in O which relate concepts to each other; $ABox$ represents all axioms which make assertions about individuals. Hereafter, in order to be compliant with the methodology for reducing more expressive languages to $\mathcal{EL}$ presented in [17, 18], we assume that $ABox$ axioms are removed with care. We use $\models$ to denote that an ontology entails an axiom (i.e. $O \models \alpha$): O^* is the deductive closure of O, i.e. the set of axioms in $\mathcal{EL}$ that can be entailed by the axioms in O.

Returning to our running example (§ 2.1), let us consider the relevant part of the deductive closure of the ISTAR ontology (O_I^*) shown in Figure 1; r_1 denotes the role `requiresOperationalCapability`.

Moreover, our ontology engineer releases a certificate assessing the *degree of confidence* for the ontology: for instance, if the engineer is requested to describe each type of operation that the US Army can perform, and he considered those for peace-keeping only, the degree of confidence of this ontology will be quite low. Formally, the ontology engineer's *degree of confidence* about the ontology is a SL opinion.

Definition 6 (Confidence Degree in the Completeness of an Ontology). *Given an ontology* O*,* $o_c^O = \langle b(c)^O, d(c)^O, u(c)^O, a(c)^O \rangle$ *is the* confidence degree in O*.*

Both the *White knowledge* and the *Black knowledge* are represented as subsets of the concepts in the ontology.

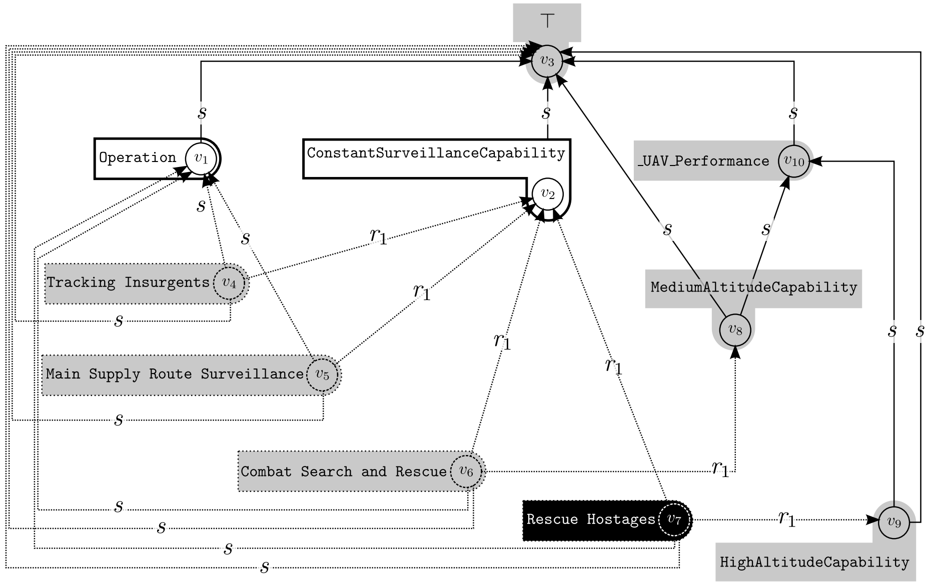
⊤
v_3
s
Operation v_1
ConstantSurveillanceCapability
v_2
_UAV_Performance v_{10}
r_1
Tracking Insurgents v_4
MediumAltitudeCapability
v_8
Main Supply Route Surveillance v_5
Combat Search and Rescue v_6
Rescue Hostages v_7
v_9
HighAltitudeCapability

Fig. 1. Part of the deductive closure of the ISTAR ontology O_I^* relevant to the scenario discussed in Section 2.1. r_1 represents the role `requiresOperationalCapability`. Figure derived following the approach described in [19].

Definition 7 (White and Black knowledge). *The* Black knowledge *(resp. the* White knowledge*) associated to the interaction between sender and receiver is* $\mathsf{B} \subseteq N_C^O$ *(resp.* $\mathsf{W} \subseteq N_C^O$*,* $\mathsf{W} \cap \mathsf{B} = \emptyset$*).*

In Figure 1, the White knowledge is v_1 and v_2 ($\mathsf{W} = \{v_1, v_2\}$), while the Black knowledge is v_7 ($\mathsf{B} = \{v_7\}$).

We introduce a machinery for evaluating the White–Black knowledge connections. For this purpose, we begin by describing a generic *semantic relationship* between two concepts representing a semantic connection, i.e. being part of a role, or of a taxonomic relation.

Definition 8 (Semantic Relationship). *A* semantic relationship $\leftrightsquigarrow \subseteq N_C^O \times N_C^O$ *is a symmetric function on the domain of concepts.*

As before, if not explicitly mentioned, we always refer to an arbitrary but fixed semantic relationship $\leftrightsquigarrow \subseteq N_C^O \times N_C^O$.

Let us introduce $\overset{R}{\leftrightsquigarrow}$, a more specific semantic relationship on the concepts names, which is parametric in the domain of the roles.

Definition 9 (Relationship $\overset{R}{\leftrightsquigarrow}$). *Given* $R \subseteq N_R^O$*,* $\overset{R}{\leftrightsquigarrow} \subseteq N_C^O \times N_C^O$ *is a symmetric relationship such that:* $v_1 \overset{R}{\leftrightsquigarrow} v_2$ *and* $v_2 \overset{R}{\leftrightsquigarrow} v_1$ *iff* $(v_1 \sqsubseteq v_2) \in O^*$ *or* $(v_1 \sqcap (\exists r.v_2)) \in O^*, r \in R$

Proposition 1. *For every* $R \in 2^{N_R^O}$*, the relationship* $\overset{R}{\leftrightsquigarrow}$ *is a semantic relationship.*

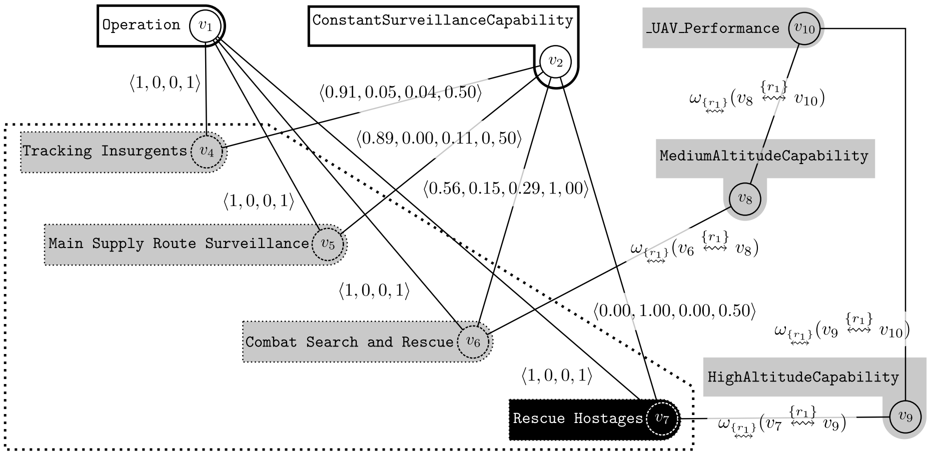
Operation v_1
ConstantSurveillanceCapability v_2
_UAV_Performance v_{10}
$\langle 1, 0, 0, 1\rangle$
$\langle 0.91, 0.05, 0.04, 0.50\rangle$
$\omega_{\{r_1\}}(v_8 \overset{\{r_1\}}{\leftrightsquigarrow} v_{10})$
Tracking Insurgents v_4
$\langle 0.89, 0.00, 0.11, 0, 50\rangle$
MediumAltitudeCapability v_8
$\langle 1, 0, 0, 1\rangle$
$\langle 0.56, 0.15, 0.29, 1, 00\rangle$
Main Supply Route Surveillance v_5
$\omega_{\{r_1\}}(v_6 \overset{\{r_1\}}{\leftrightsquigarrow} v_8)$
$\langle 1, 0, 0, 1\rangle$
$\langle 0.00, 1.00, 0.00, 0.50\rangle$
$\omega_{\{r_1\}}(v_9 \overset{\{r_1\}}{\leftrightsquigarrow} v_{10})$
Combat Search and Rescue v_6
HighAltitudeCapability
$\langle 1, 0, 0, 1\rangle$
Rescue Hostages v_7
$\omega_{\{r_1\}}(v_7 \overset{\{r_1\}}{\leftrightsquigarrow} v_9)$
v_9

Fig. 2. Semantic relationship graph induced on O_I by the semantic relation $\overset{\{r_1\}}{\leftrightsquigarrow}$. The depicted SL labelling have the following meaning: *receiver* knows for sure the taxonomic relations between v_1 and it subclasses v_4, v_5, v_6 and v_7; there is a large amount of evidence for believing that *receiver* knows the relationships $v_4 \overset{\{r_1\}}{\leftrightsquigarrow} v_2$ and $v_5 \overset{\{r_1\}}{\leftrightsquigarrow} v_2$; the set of evidence is smaller in the case of v_6, but *sender* is very biased; *sender* has no clue about $v_7 \overset{\{r_1\}}{\leftrightsquigarrow} v_2$. The other SL labelling do not affect our running example and are left in their symbolic form.

A semantic relationship induces an indirect graph (*semantic relationship graph*). We then label each edge in such a graph with a SL opinion, which represents the likelihood that the *receiver* is aware of the semantic relationship represented by such an edge. These SL opinions can be derived from past interactions with the *receiver* [20], or by *a priori* knowledge.

Definition 10 (SL Labelling of Semantic Relationship). *The function* $\omega_{\leftrightsquigarrow} : \leftrightsquigarrow \mapsto \mathbb{O}$ *is a* SL labelling of $\leftrightsquigarrow$*:* $\omega_{\leftrightsquigarrow}(v_1 \leftrightsquigarrow v_2) = \omega_{\leftrightsquigarrow}(v_2 \leftrightsquigarrow v_1)$.

In our running example, the *semantic relationship graph* induced by the semantic relation $\overset{\{r_1\}}{\leftrightsquigarrow}$ on O_I is depicted in Figure 2. The meaning of the area comprised in the dotted lines will be explained following Alg. 1.

We can now define a *semantic path* as a path in the semantic relationship graph.

Definition 11 (Semantic Path). *A* semantic path between v_A and v_B $p_{\leftrightsquigarrow}(v_A, v_B)$ *is a sequence of nodes* $p_{\leftrightsquigarrow}(v_A, v_B) = \langle v_1, \ldots, v_n\rangle$ *s.t.* $v_1 = v_A$, $v_n = v_B$, *and* $\forall i < n$, $v_i \leftrightsquigarrow v_{i+1}$ *holds.*

In order to assess the likelihood of surmising a concept, we rely on the intuition that the closer — in terms of semantic relationship — two concepts, the greater the likelihood to predict one from the other. Therefore, we need a measure of the *semantic distance w.r.t. a semantic relationship* between two sets of concepts. To this end, we first need to define the *set of minimal semantic paths* between two set of concepts, which is necessary to assess the distance among them.

Definition 12 (Set of Minimal Paths). *Given* $Z_1, Z_2 \subseteq N_C^O$*, the* set of minimal paths between Z_1 and Z_2 *is*

$$P_{\leftrightsquigarrow}(Z_1, Z_2) = \{p_{\leftrightsquigarrow}(v_1, v_2) \mid v_1 \in Z_1, v_2 \in Z_2, \\ p_{\leftrightsquigarrow}(v_1, v_2) \in \underset{p_{\leftrightsquigarrow}(v_1, v_2)}{\arg\min} (|p_{\leftrightsquigarrow}(v_1, v_2)|)\}$$

such that $\forall v_1 \in Z_1, v_2 \in Z_2$, $\exists! p_{\leftrightsquigarrow}(v_1, v_2) \in P_{\leftrightsquigarrow}(Z_1, Z_2)$.

Then, the semantic distance w.r.t. a semantic relationship between two sets of concepts is the maximum among the lengths of minimal paths between them.

Definition 13 (Semantic distance). *Given* $Z_1, Z_2 \subseteq N_C^O$*, the* semantic distance w.r.t. the semantic relationship $\leftrightsquigarrow$ between Z_1 and Z_2 *is*

$$d_{\leftrightsquigarrow}(Z_1, Z_2) = \max_{p_{\leftrightsquigarrow}(v_1, v_2) \in P_{\leftrightsquigarrow}(Z_1, Z_2)} (|p_{\leftrightsquigarrow}(v_1, v_2)|)$$

Note that $Z_1 \subseteq N_C^O$, $d_{\leftrightsquigarrow}(Z_1, Z_1) = 0$.

Finally, in order to take into account the *receiver*'s believed knowledge, a *cumulative SL labelling* between two concepts is defined as the conjunction of the SL labelling of the minimal paths between them.

Definition 14 (Cumulative SL Labelling). *Given* $\omega_{\leftrightsquigarrow}$ *a SL labelling of* $\leftrightsquigarrow$*,* $Z_1, Z_2 \subseteq N_C^O$*, the* cumulative SL labelling between Z_1 and Z_2 *is*

$$\zeta_{Z_1 \leftrightsquigarrow Z_2} = \bigwedge_{p_{\leftrightsquigarrow}(v_A, v_B) \in P_{\leftrightsquigarrow}(Z_1, Z_2)} \; \bigwedge_{i < |p_{\leftrightsquigarrow}(v_A, v_B)|} \omega_{\leftrightsquigarrow}(v_i \leftrightsquigarrow v_{i+1})$$

Although multiple minimal paths can exist between $v_1 \in Z_1$ and $v_2 \in Z_2$, we require (Def. 12) that only one of them is included in $P_{\leftrightsquigarrow}(Z_1, Z_2)$. We do not enforce a specific method for choosing the one: the most reasonable is to include the minimal path which maximize a metric, like the cumulative SL labelling (Def. 14).

4 SOF: a Framework for Ontology Obfuscation

We are now able to provide a formalization for *semantic obfuscation*. We call our approach the *Semantic Obfuscation Framework* (SOF). This measures the *quality of obfuscation* in terms of the "likelihood" that the *receiver* will surmise the Black knowledge from the White knowledge. First, we formalize the concept of "surmise" based on a semantic relationship.

Definition 15 (Semantic Surmise). *The* surmise from a node v_1 *is* $S_O^{\leftrightsquigarrow}(v_1) = \{v_1\} \cup \{v_2 \mid v_1 \leftrightsquigarrow v_2\}$.

With a little abuse of notation, we define surmise on a set of concepts.

Definition 16 (Semantic Surmise of a Set). *Given $Z_1 \subseteq N_C^O$:*

$$S_O^{\leftrightsquigarrow}(Z_1) = Z_1 \cup \bigcap_{v_1 \in Z_1} S_O^{\leftrightsquigarrow}(v_1)$$

Moreover, given the confidence degree, we define a function for estimating the dimension (in terms of number of concept) of the "perfect" domain ontology.

Definition 17 (Estimative Function). *A function $\phi : \mathbb{O} \times \mathbb{R} \mapsto \mathbb{R}$ is an* estimative function *iff $\phi(o,n) \geq n$ and $\phi(o,n) = n$ iff $o = \langle 1, 0, 0, \cdot \rangle$; and $\phi(o,n+m) = \phi(o,n) + \phi(o,m)$.*

In particular, let us introduce a *cautious estimative function*.

Definition 18 (Cautious Estimative Function). *The* cautious estimative function ϕ^C : $\mathbb{O} \times \mathbb{R} \mapsto \mathbb{R}$ *is defined as: $\phi^C(o,n) = \dfrac{n}{\mathbb{E}[o]}$. If $\mathbb{E}[o] = 0$, $\phi^C(o,n) = K \gg 1$.*

Proposition 2. *The* cautious estimative function ϕ^C *is an estimative function.*

We can provide now a computational procedure (Alg. 1) for deriving the degree of likelihood that a *receiver* will surmise the Black knowledge from the White knowledge. Alg. 1 requires as input an ontology O, the White and Black knowledge (resp. W and B), a semantic relationship $\leftrightsquigarrow$, an estimative function ϕ, and the confidence degree o_c^O.

Algorithm 1 performs several computations. First of all, it determines S (l. 1 of Alg. 1), namely the minimum set of surmised concept names from the White Knowledge which includes the Black Knowledge also.

It then considers the focused frame of discernment composed by two disjoint primitive states, viz. B and $\mathsf{S} \setminus \mathsf{B}$. In particular, it uses the cumulative labelling between W and both B and $\mathsf{S} \setminus \mathsf{B}$ (l. 2 of Alg. 1) for computing the mass assignment of the two primitive states (l. 3 of Alg. 1). To this end, Alg. 1 exploits the well-known relationship between SL opinions on a focused frame of discernment and the beta distribution [7].

In order to prove the soundness of such approach, in the case of perfect knowledge, the SL opinion must collapse to a traditional probability value. The following propositions shows that this is the case.

Proposition 3. *If $o_c^O = \langle 1, 0, 0, 1 \rangle$ and $\forall v_1, v_2 \in N_C^O, \omega_{\leftrightsquigarrow}(v_1 \leftrightsquigarrow v_2) = \langle 1, 0, 0, 1 \rangle$, o_B is equivalent to considering* B *as the set of outcomes of an experiment on the sample space* S*, and thus computing the probability of* B.

Proof. *From Def. 17, $\phi(o_c^O, n) = n$; from $\omega_{\leftrightsquigarrow}(v_1 \leftrightsquigarrow v_2) = \langle 1, 0, 0, 1 \rangle$, $\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}} = \langle 1, 0, 0, 1 \rangle$ and $\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{S} \setminus \mathsf{B}} = \langle 1, 0, 0, 1 \rangle$. Therefore, $r = |\mathsf{B}|$, $s = |\mathsf{S} \setminus \mathsf{B}|$, $V = 0$.* □

In our running example, $S_I = \{v_4, v_5, v_6, v_7\}$ — the area comprised by dotted line in Figure 2 — is the minimum set of surmised concepts needed to reach the Black knowledge from the white knowledge. In the case of perfect knowledge, the probability to surmise the Black knowledge is 0.25 (cf. Proposition 3).

Let us consider instead the case where there is complete (unbiased) uncertainty about the completeness of O_I, i.e. $o_c^{O_I} = \langle 0, 0, 1, 0.5 \rangle$. Therefore, $\phi^C(o_c^{O_I}, n) = 2 \cdot n$. Let us also suppose that the SL labelling are as depicted in Figure 2.

Algorithm 1 Procedure for deriving the degree of likelihood about surmising Black knowledge

```
$\texttt{DegreeBlackSurmising}(O, \mathsf{W}, \mathsf{B}, \leftrightsquigarrow, \phi, o_c^O)$
  compute $\mathsf{S} = \min \left( \bigcup_{Z_1 \subseteq N_C^O,\ d_{\leftrightsquigarrow}(\mathsf{W}, Z_1) \leq d_{\leftrightsquigarrow}(\mathsf{W},\mathsf{B})} S_O^{\leftrightsquigarrow}(Z_1) \right)$ s.t. $\mathsf{B} \subseteq \mathsf{S}$
  compute the cumulative labelling $\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}} = \langle b(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}}), d(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}}), u(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}}), a(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}}) \rangle$,
    and $\zeta_{\mathsf{W} \leftrightsquigarrow (\mathsf{S} \setminus \mathsf{B})} = \langle b(\zeta_{\mathsf{W} \leftrightsquigarrow (\mathsf{S} \setminus \mathsf{B})}), d(\zeta_{\mathsf{W} \leftrightsquigarrow (\mathsf{S} \setminus \mathsf{B})}), u(\zeta_{\mathsf{W} \leftrightsquigarrow (\mathsf{S} \setminus \mathsf{B})}), a(\zeta_{\mathsf{W} \leftrightsquigarrow (\mathsf{S} \setminus \mathsf{B})}) \rangle$
  compute
    $b(B) = \frac{r}{V + \phi(o_c^O, r + s)}$        $d(B) = \frac{s}{V + \phi(o_c^O, r + s)}$
    $u(B) = \frac{V + \phi(o_c^O, r + s) - r - s}{V + \phi(o_c^O, r + s)}$        $a(B) = \min\{a(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}}), a(\zeta_{\mathsf{W} \leftrightsquigarrow (\mathsf{S} \setminus \mathsf{B})})\}$
  where
    $r = \begin{cases} \frac{|\mathsf{B}| \cdot b(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}})}{u(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}})} & u(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}}) \neq 0 \\ |\mathsf{B}| & \text{otherwise} \end{cases}$        $s = \begin{cases} \frac{|\mathsf{S} \setminus \mathsf{B}| \cdot b(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{S} \setminus \mathsf{B}})}{u(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{S} \setminus \mathsf{B}})} & u(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{S} \setminus \mathsf{B}}) \neq 0 \\ |\mathsf{S} \setminus \mathsf{B}| & \text{otherwise} \end{cases}$
    $V = u(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{S} \setminus \mathsf{B}}) \cdot u(\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}}) \cdot |\mathsf{S}|$
  **return** $\langle b(B), d(B), u(B), a(B) \rangle$
```

Therefore, $\zeta_{\mathsf{W} \leftrightsquigarrow \mathsf{B}} = \omega_{\underset{\leftrightsquigarrow}{\{r_1\}}}(v_7 \overset{\{r_1\}}{\leftrightsquigarrow} v_2)$ and $\zeta_{\mathsf{W} \leftrightsquigarrow (\mathsf{S}_I \setminus \mathsf{B})} = \langle 0.61,\ 0.36, 0.03, 0.25 \rangle$. The result of the computation is $o_B = \langle 0, 00,\ 0.49,\ 0.51,\ 0.25 \rangle$: $\mathbb{E}[o_B] = 0.12$.

Intelligence thus informs the CJTF Commander that the coalition local partners will surmise that the operation will be a hostage rescue given the request for support for constant surveillance of the POI with a 12% probability. Since the CJTF considers 12% a reasonable risk, he requests for constant surveillance of the member of the VEO Sumer. This leads the coalition forces to the location of the two American diplomats and to the solution of the situation.

5 Discussion and Conclusions

In this paper we introduce *SOF*, a Semantic Obfuscation Framework, which includes a sound and effective procedure for evaluating the likelihood that the *receiver* of some pieces of information will derive some additional information that the *sender* desires to keep private. Related works can be found in two different areas. First, which is the main line of investigation that motivates this work, regards multi-agents strategic interactions and the assessment of risk-benefit trade-off of information sharing [1,2,4,5,6,20,21]. In this context, past approaches relied on information theoretic metrics and on quantitative data. We claim that a qualitative representation of the domain like the one proposed in this paper has several advantages. First, encompassing contextual knowledge — e.g. the *receiver* knows that the American diplomats have been kidnapped by reading the newspaper — is possible by enlarging the ontology and applying the same methodology. Second, such a formalization improves the users' awareness of inferences that could be surmised. For instance half of the elements in S_I *requires operational capability* related to Unmanned aerial vehicles (UAVs) (cf. Figure 2): this might suggest that there could be a semantic connection between `ConstantSurveillanceCapability` and `_UAV_Performance`. In future work we will enlarge this discussion and provide both a formal and an experimental comparison with the quantitative approaches.

The second trend in the literature regards ontology mapping under uncertainty. Among others, the relevant work utilizes either supervised machine learning [22], or unsupervised machine learning using Bayesian networks [23], Markov networks [24], and information theory [25]; or Fuzzy logic approaches [26]. In addition, [27, 28] address the topic of uncertainty management using qualitative techniques, in particular beliefs networks and argumentation. However, the present paper is the first approach aimed at "obfuscating" shared pieces of knowledge for a specific purpose: i.e., to prevent the revelation of Black knowledge. This shift of paradigm thus makes a comparison with ontology mapping approaches difficult.

Although this paper is focused on sharing set of concepts, a natural evolution would involve sharing set of axioms. In this area several approaches have been proposed for modifying an ontology to conceal sensitive information both in $\mathcal{EL}$ [29] and in more expressive languages [30, 31]. A discussion on this topic is beyond the scope of this paper and is left for future work.

Finally, it is worth mentioning that the impact of this line of research can spread beyond the military context, which we considered in this paper: for instance it could form the basis for an investigation on how to improve users' awareness of the confidentiality of the information shared across social media.

On the other hand, considering the point of view of the receiver, such approach can identify gaps in the received information, a relevant topic within intelligence analysis. In intelligence analysis often the key information is lacking, and analysts must factor the impact of missing data in their confidence when judging a situation. Failing to recognize absence of evidence is one of the most important causes of cognitive bias [32]. However, it is hard to identify what is known to be missing and why (e.g. different clearance levels, hidden agenda, ...). Reversing our approach can support analysts in the quest for further evidence to fill the gaps required to deliver more effective intelligence.

Acknowledgments

The authors thank the anonymous reviewers for their helpful comments.

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

References

1. Chakraborty, S., Bitouz, N., Srivastava, M.B., Dolecek, L., Bitouze, N., DolecekLara: Protecting Data Against Unwanted Inferences. In: 2013 Information Theory Workshop. (2013)
2. Chakraborty, S., Raghavan, K.R., Johnson, M.P., Srivastava, M.B.: A framework for context-aware privacy of sensor data on mobile systems. In: Proceedings of the 14th Workshop on Mobile Computing Systems and Applications - HotMobile '13, New York, New York, USA, ACM Press (February 2013) 1
3. Goffman, E.: Strategic Interaction. Basil Blackwell Oxford (1970)
4. Cerutti, F., Chakraborty, S., de Mel, G., Gibson, C., Kaplan, L., Oren, N., Pipes, S., Sensoy, M., Srivastava, M., Zafer, M.: A Principled Framework for Risk vs. Utility Management in Coalition Networks. In: Annual Fall Meeting of ITA. (2013)
5. Bisdikian, C., Tang, Y., Cerutti, F., Oren, N.: A Framework for Using Trust to Assess Risk in Information Sharing. In Chesñevar, C., Onaindia, E., Ossowski, S., Vouros, G., eds.: Agreement Technologies – Second International Conference, AT 2013 Beijing, China, August 2013 Proceedings, Springer Berlin Heidelberg (2013) 135–149
6. Pipes, S., Hardill, B., Gibson, C., Srivastava, M., Bisdikian, C.: Exploitation of Distributed, Uncertain and Obfuscated Information Experimental application of trust and obfuscation techniques in distributed middleware. Technical report (2012)
7. Jøsang, A.: A Logic for Uncertain Probabilities. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems **9**(3) (2001) 279–311
8. Şensoy, M., Fokoue, A., Pan, J.Z., Norman, T.J., Tang, Y., Oren, N., Sycara, K.: Reasoning about uncertain information and conflict resolution through trust revision. In: Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems, International Foundation for Autonomous Agents and Multiagent Systems (May 2013) 837–844
9. Cerutti, F., Toniolo, A., Oren, N., Norman, T.J.: An Empirical Evaluation of Geometric Subjective Logic Operators. In Chesñevar, C., Onaindia, E., Ossowski, S., Vouros, G., eds.: Agreement Technologies – Second International Conference, AT 2013 Beijing, China, August 2013 Proceedings, Springer Berlin Heidelberg (2013) 90–104
10. Jøsang, A., Pope, S., Marsh, S.: Exploring Different Types of Trust Propagation. In: Proceedings of the 4th International Conference on Trust Management (iTrust'06). (2006)
11. Sullivan, P.: Scenario stabilization operations. `https://www.usukitacs.com/sites/default/files/Scenario%20stabilization%20operations%20v2.doc` (2014)
12. Şensoy, M., de Mel, G., Vasconcelos, W.W., Norman, T.J.: Ontological logic programming. In: Proceedings of the International Conference on Web Intelligence, Mining and Semantics - WIMS '11, New York, New York, USA, ACM Press (May 2011) 1
13. Şensoy, M., Vasconcelos, W.W., Norman, T.J.: Combining Semantic Web and Logic Programming for Agent Reasoning. In Dechesne, F., Hattori, H., Mors, A., Such, J., Weyns, D., Dignum, F., eds.: Advanced Agent Technology. Volume 7068 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2012) 428–441
14. Dempster, A.P.: A Generalization of Bayesian Inference. Journal of the Royal Statistical Society. Series B (Methodological) **30**(2) (1968) pp. 205–247
15. Shafer, G.: A Mathematical Theory of Evidence. Princeton University Press (1976)
16. Baader, F.: Terminological Cycles in a Description Logic with Existential Restrictions. In Gottlob, G., Walsh, T., eds.: IJCAI, Morgan Kaufmann (2003) 325–330
17. Ren, Y., Pan, J.Z., Zhao, Y.: Towards Soundness Preserving Approximation for ABox Reasoning of OWL2. In Haarslev, V., Toman, D., Weddell, G.E., eds.: Proceedings of the 23rd International Workshop on Description Logics (DL 2010), Waterloo, Ontario, Canada, May 4-7, 2010 2010), Waterloo, Ontario, Canada, May 4-7, 2010. (2010)

18. Ren, Y., Pan, J.Z., Zhao, Y.: Soundness Preserving Approximation for TBox Reasoning. In Fox, M., Poole, D., eds.: Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010. (2010) 351–356
19. Mitra, P., Wiederhold, G., Kersten, M.: A Graph-Oriented Model for Articulation of Ontology Interdependencies. In Zaniolo, C., Lockemann, P., Scholl, M., Grust, T., eds.: Advances in Database Technology EDBT 2000. Volume 1777 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2000) 86–100
20. Pipes, S., Cerutti, F., Chakraborty, S., Gibson, C., Norman, T.J., Oren, N., Srivastava, M.: Architecture Options for Realization of Inference Management in Information Fabric. `https://www.usukitacs.com/node/2633` (2014)
21. Bisdikian, C., Tang, Y., Cerutti, F., Oren, N.: Reasoning about the Impacts of Information Sharing. Information Systems Frontiers Journal, under submission
22. Duan, S., Fokoue, A., Srinivas, K.: One Size Does Not Fit All: Customizing Ontology Alignment Using User Feedback. In: The Semantic Web ISWC 2010. (2010) 177–192
23. Mitra, P., Noy, N., Jaiswal, A.R.: Ontology mapping discovery with uncertainty. In: Proc. 4th International Semantic Web Conference (ISWC). Volume 3729. (2005) 537–547
24. Albagli, S., Ben-Eliyahu-Zohary, R., Shimony, S.E.: Markov network based ontology matching. Journal of Computer and System Sciences **78**(1) (2012) 105–118
25. Blasch, E.P., Dorion, É., Valin, P., Bossé, E.: Ontology alignment using relative entropy for semantic uncertainty analysis. In: Aerospace and Electronics Conference (NAECON), Proceedings of the IEEE 2010 National, IEEE (2010) 140–148
26. Gal, A.: Managing uncertainty in schema matching with top-k schema mappings. In: Journal on Data Semantics VI. Springer (2006) 90–114
27. Nagy, M., Vargas-Vera, M., Motta, E.: Dssim-managing uncertainty on the semantic web. (2007)
28. Trojahn, C., Quaresma, P., Vieira, R.: An extended value-based argumentation framework for ontology mapping with confidence degrees. In: Argumentation in Multi-Agent Systems. Springer (2008) 132–144
29. Konev, B., Walther, D., Wolter, F.: Forgetting and uniform interpolation in large-scale description logic terminologies. In: 21st International Joint Conference on Artificial Intelligence, Morgan Kaufmann Publishers Inc. (July 2009) 830–835
30. Koopmann, P., Schmidt, R.: Forgetting Concept and Role Symbols in $\mathcal{ALCH}$-Ontologies. In McMillan, K., Middeldorp, A., Voronkov, A., eds.: Logic for Programming, Artificial Intelligence, and Reasoning. Volume 8312 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2013) 552–567
31. Wang, K., Wang, Z., Topor, R., Pan, J., Antoniou, G.: Concept and Role Forgetting in $\mathcal{ALC}$ Ontologies. In Bernstein, A., Karger, D., Heath, T., Feigenbaum, L., Maynard, D., Motta, E., Thirunarayan, K., eds.: The Semantic Web - ISWC 2009. Volume 5823 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2009) 666–681
32. Heuer, R.: Psychology of intelligence analysis. US Government Printing Office (1999)