Context-Enhanced Information Fusion: Applications

Lauro Snidaro
Dept. Mathematics, Computer Science and Physics
University of Udine
ITALY

lauro.snidaro@uniud.it

ABSTRACT

In these notes, three main application lines for context in Information Fusion will be presented. In the first, context-aware target tracking will be briefly introduced with references to available works in the literature. The more extended second part will discuss the role of context in inference for Situation Assessment in a maritime scenario. The framework of Markov Logic Networks will be applied for fusing relational information coming from heterogeneous sources. In the third part, a de-contextualization case will be discussed for military intelligence reasoning.

1.0 CONTEXT FOR TARGET TRACKING

This section briefly discusses how context can be applied to target tracking. Since a good number of context-aware tracking solutions exist in the literature [1], [2] and given the limited space available here, the topic will be here only introduced by briefly discussing a possible architecture [3]. The Resource Management Module (RMM) presented is intended to be part of an engine for multi-sensor fusion which is able to combine observations from multiple and possibly heterogeneous sensors for tracking and classification. These two tasks are the main duties of the engine supported by a contextual database that encodes a priori knowledge on the observed environment (e.g. map of the area, weather map, etc.). The RMM would be in charge of assigning a quality measure to the measurements generated by the sensors considering possible contextual effects that may hamper their performance. The architecture is deliberately very simple for sake of introducing here the topic. The already cited references above provide pointers to algorithmic solutions available in the literature. Additional details on the architecture presented here can be found in [3]. A more in depth and general discussion on context-aware fusion architectures can be found in [4].

The overall RMM processing chain is shown in Figure 1. The module receives observations from sensors and contextual information as inputs. A first selection step is performed at this point by weighting or discarding sensor measurements that are likely to be false alarms (e.g. noise, reflections, etc.). This is particularly true for tracking, since position measurements can be easily checked against spatial and structural constraints. Observations validated by context are then evaluated for the information gain they could provide to the estimate computed by the engine at the previous time instant regarding a given target. According to their informative value, a subset of all the measurements is selected and passed to the fusion engine that will generate the new estimate. This loop is repeated for each collection of new observations received by the sensors.

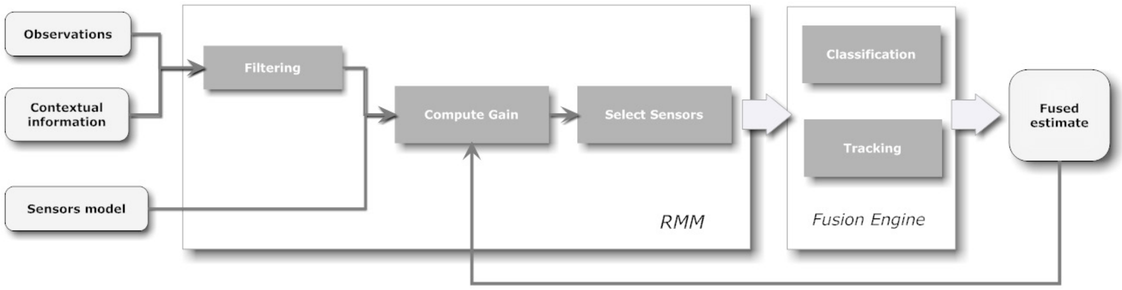

Figure 1: Data Flow Through Modules. The filtering step is performed through the exploitation of contextual information into the RMM.

The RMM will consider the problem of assessing to which extent the measurements coming from a sensor should be used in the fusion process by modelling the associated information gain. This second step is here included to model a possible calculation of the *relevance* [5] of a given observation for the fusion process.

In an ideal framework, this applies to several tasks as, for instance, Tracking and Classification. With this respect, even if couched in the same framework, tracking and classification can be treated separately even if readings are coming from the same platform. For instance, there can be a sensor that provides very good localization with a poor associated classification result, and it can be used for tracking but not for classification. Conversely, one source may provide good data for recognition, but at the same time being affected by large localization errors. In this case, the sensor output can be used for classification but not for localization. In the following, we will discuss the how the filtering step could be performed for the tracking process by taking advantage of available contextual information.

Considering the available scenario information, a pre-filtering phase can be used to remove very uncertain observations, considering variables as weather or time of the day. Imagine, for instance, the case of a target moving along a city street and suppose that we want to estimate its state x as the vector of Cartesian bidimensional coordinates. Suppose now that the observation $z(t)$ at time t is checked against an urban map of the monitored area resolving $z(t+1)$ as falling inside a building. Now, given the fact that we know that the sensor has no see-through-walls capability, this could be explained as an occasional quirk of the sensor and could be easily filtered out by the tracking algorithm (e.g. Kalman filter, particle filter, etc.). Especially if $z(t)$ is resolved inside a building while both the previous state $x(t-1)$ and the next measurement $z(t+1)$ do not.

Unfortunately, in real-world monitoring applications it often happens that a sensor provides a sequence of unreliable observations due to partial occlusion of the target, unfavourable weather conditions, sun blinding, persistent reflections, etc. In these cases, tracking can be severely disrupted providing an unreliable estimate of the target's position and trajectory.

The pre-filtering step exploits, therefore, different and diversified contextual information as a means to filter observations, as shown in Figure 2. The aim of integrating contextual information into tracking systems is to better refine and optimize the task according to the observations provided by the sensors and to prior high-level knowledge of the environment, which is coded as context. Contextual information can be a key factor in determining the state of an entity of interest, as it can dramatically impact on the reliability of an observation.

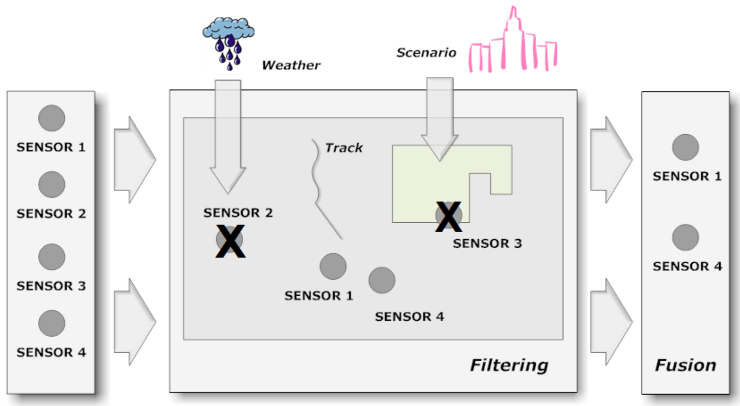

Figure 2: Measurements Received from Multiple Sensors are Filtered According to Contextual Information.

Checking the measurements against a map of the monitored area is a form of contextual knowledge inclusion that could, as in the latter example, provide an insight on the reliability of the sensor in a specific situation. Knowing the sensing capabilities of a sensor is another form of contextual knowledge that could be exploited conveniently. In the previous example, knowing that the sensor has no see-through-walls capability *and* the fact that the last few measurements fall inside a building can help us in concluding that those measurements may be affected by a form of bias and thus be unreliable. The sensor may be in fact persistently experiencing one or more of the disturbing conditions mentioned above.

To be strict, one can discard the sensors that give measurements not compatible with the reliability assigned by contextual information, adopting thus a *pruning* strategy. Alternatively, instead of getting rid of sensors observations, the measurements can be combined by weighting them (*discounting*) with respect to the reliability factor given by context analysis.

For more details, the reader is referred to surveys of context-aware tracking algorithms can be found in [1], [2], while [7] can serve as a tutorial on the topic.

2.0 CONTEXT FOR SITUATION ASSESSMENT

The Situation Assessment (SA) process in Information Fusion (IF) security systems has a clear goal: building and updating a situational picture of the scenario under consideration. In the maritime domain, the scenario is generally very dynamic in time and comprises a large number of entities and actors operating in a complex environment. SA aims at explaining the observed events (mainly) by establishing the entities and actors involved, inferring their goals, understanding the relations existing (permanently or temporarily) between them, the surrounding environment, and past and present events. It is therefore evident how the SA process inherently hinges on understanding and reasoning about relations. SA processes are particularly complex and critical for large-scale scenarios such as those related to border and port security, where suspicious activities need to be detected as the needle in a haystack of largely predominant “pattern of life” activities.

Here, we are particularly interested in representing and reasoning about relations and events. This would allow the system to be able to capture the relationships existing between elements of the scenario via the recognition of sequences of events (complex events) that encode situations of interest with possible dangerous/disastrous outcomes.

We will here provide a summary on our ongoing work [8] on applying methods belonging to the rapidly growing area of Statistical Relational Learning to the task of SA in information fusion systems. Specifically, we will discuss how within the framework of Markov Logic Networks (MLNs) different types of complex relational knowledge (e.g. contextual) with associated uncertainty can be fused together along with observational data for event understanding

2.1 State-of-the-Art

State-of-the-art situation assessment (SA) systems are able to deal with vast amounts of data and information also of a heterogeneous kind. Their goal is to provide a constantly updated situational picture about the observed environment or set of entities to an operator in order to facilitate human decision making. Updating the current system representation of the situation is generally performed by acquiring, through sensors or other sources of information, new observations which provide a possibly incomplete and uncertain view.

Currently, low-level sensory data is the main source of information used to understand the observed evolving scenario and to identify anomalous conditions; in particular, up to now maritime surveillance heavily relies on the Automatic Identification System (AIS), coastal radars, space-based imagery, and other sensors, to form a picture in which the operator can recognize complex patterns and make decisions.

The common thread that unites many works in the literature (see [8] for a recent literature review) is the definition of an expert system, that aims at detecting a set of anomalous behaviours or potential threats. Subject matter experts define a knowledge base (KB), which comprises the possible abnormal patterns the target could follow; then, on the top of it, a reasoning engine queries the occurrence of an anomaly for a target object in an arbitrary time instant.

The reasoner is usually fired by low-level observations provided by sensors, covering in this way the majority of abnormal situations in the domain; however, it is interesting to notice how anomalous behaviours do not always follow standard trends or well-known patterns, especially if related solely to vessels movements, but sometimes they take the form of seemingly unrelated activities on a larger scale [9]. Ship-centric focus should be replaced by a broader vision, where the ideal situational awareness system should then be flexible and adaptive enough to integrate both low-level and high-level information (Figure 3), detecting anomalous or suspicious conditions by reasoning on manifest or uncertain data, but also on (apparently irrelevant) relations among objects, which may reveal unobserved coincidences. This requires the interplay of both deductive and abductive inferencing processes. How this involvement can be obtained, on both theoretical and applicative levels, is a crucial point, and is subject of ongoing research [13], [10].

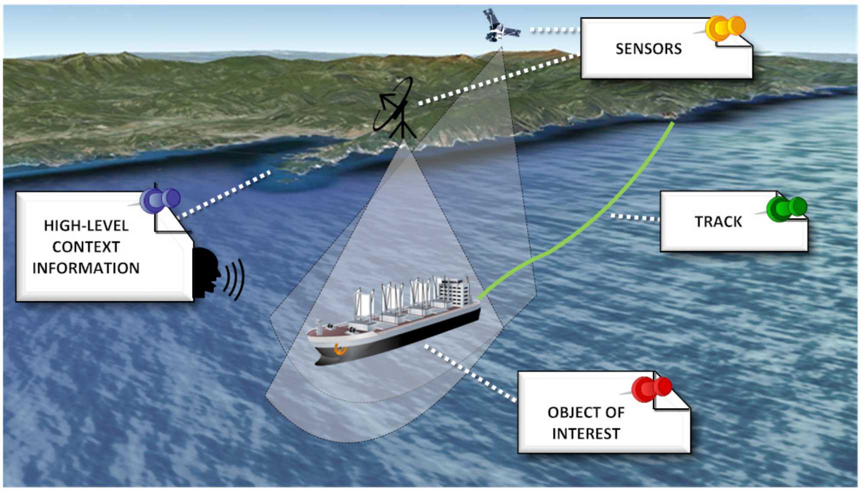

Figure 3: In Situation Awareness for Maritime Domain. Sensory data must be coupled with high-level information and contextual data.

The maritime domain is a daunting scenario for testing such systems, because of many factors:

- Its challenging nature where the coverage of wide areas is given by discontinuous and intermittent sensory data;
- Its well-known commercial policies and practices which can suggest normalcy behaviour patterns;
- The presence of local contextual information, stable in time, which can depict alternative indicators of multi-layered situations; and
- The urgency for systems capable to provide effective and advanced warning to promote countermeasures to illicit activities.

The integration of contextual knowledge can greatly enhance the performance of an SA system [8]. Despite its value, the representation and use of context is often poorly integrated, if not absent, even if the richness and completeness of this information is extremely useful to properly interpret the available stream of raw sensor data from a multitude of points of view (security, safety, economical or environmental situation, etc.).

Qualitative high-level knowledge can help to infer about hidden states from low-level data generated by sensors, other fusion processes or human reports. In other words, context is a powerful means to picture a broader and deeper operational situation, as it can reduce uncertainties where normally analysts would need to be consulted.

In these notes, we show how MLNs can be exploited to encode uncertain knowledge, fuse data coming from multiple (and possibly heterogeneous) sources, and perform reasoning on incomplete data. One key point of using the MLNs for reasoning is their ability to reason with incomplete or missing evidence, which is a crucial feature hardly found in other approaches, but sought after especially in the maritime domain, where the data is often inaccurate, delayed or simply not available. Another advantage with respect to other systems, is the fact that MLNs support inconsistencies or contradictions in the knowledge base, which is a problem when different experts provide contributes to it. This avoids non-trivial knowledge engineering techniques to be performed in order to guarantee rules consistency.

Here we use Markov Logic Networks (MLNs) [11] to detect two possible anomalous conditions in maritime domain, a rendezvous at sea and a hazardous combination of cargo ships in a harbour.

We use exemplary scenarios, the first one derived from experts' suggestions gathered at the NATO STO Centre for Maritime Research and Experimentation and the second one expanded from [12], to highlight how unobserved complex events could be built by logical combination of simpler evidence, and how contextual information is extremely valuable in many conditions.

MLNs present advantages suited to our domain as they:

- Support reasoning with missing or partial observations (incomplete evidence), they allow to encode expert rules and relational knowledge with an associated degree of uncertainty;
- Are able to encode the relational knowledge among objects and entities in the scenario; and
- Are able to handle contradictions and inconsistencies.

In the following, we would like to show how to tap via MLNs the expressive power of first-order logic (FOL) and the probabilistic uncertainty management of Markov networks in order to detect anomalies via reasoning on uncertain knowledge. Specifically, the following points will be discussed:

- Clarifying the concepts of event (simple and complex) and anomaly in the scope of fusion terminology;
- Explicitly explaining how simple and complex events can be encoded in the form of FOL formulas with associated degree of uncertainty in maritime domain;

- Demonstrating how MLNs could provide a powerful tool for fusing heterogeneous sources (e.g. a priori, contextual, sensory, etc.) of information for situation assessment by being able to express unobserved complex events by logical combination of simpler evidences; and
- Developing a mechanism to evaluate the level of completion of complex events as this calculation is not directly solvable within the MLNs framework.

2.2 Events and Anomalies

Events and anomalies can be considered fundamental building blocks for developing an informed situational picture of the environment. In this section, we briefly provide the necessary definitions of these concepts. Further details and discussion can be found in [8].

For our purposes, an event is a "significant occurrence or happening". It can be subdivided in *simple*, when we consider the variation of a quantity or state, or *complex*, which is a combination of atomic or complex activities.

An *anomaly* can be considered a critical event to which the system is generally called to react to. Anomalies can be simple events (like in the case of an observable quantity, e.g. oil pressure, measured by an instrument that passes a critical threshold) or complex ones (e.g. hydraulics failure). In the perspective of preventing anomalies, the SA system might require some form of reasoning to be performed in order to anticipate possible dangerous events before they happen.

Following the description given in the above sections and taking into account JDL levels [13], our position here is that whenever the system detects any appreciable variation of input data of any level, a corresponding event is generated. Table 1 shows some examples of events generated from data at different levels. For instance, the detection of presence or absence of AIS signal is something that can be considered at the bottom of the JDL hierarchy, while the speed of a vessel is a feature of its state and belongs to Level 1. Two stopped vessels very close out at sea is a relation between two entities and helps defining the current situation (JDL level 2). It is not true then, that, to flag a situation as anomalous, data and information have to bubble up through the levels following increasing processing and refinement steps. Anomalies can be generated from data of every kind and level as shown in Table 1. For example, the absence of AIS signal can be directly considered something anomalous, as well as a speeding boat or a rendezvous out at sea. Also, anomalies could be generated both from simple and complex events.

Table 1: Examples of Events and Anomalies at Different JDL Levels [8].

Level	Event	Type	Anomaly
0	Absence of AIS signal	Simple	AIS off
1	Vessel increased speed	Simple	Vessel over speed limit
2	Vessel X stopped, Vessel Y stopped, Vessel X and Y are close	Complex	Rendezvous

2.3 Markov Logic Networks

The probabilistic reasoning framework of Bayesian networks, despite their widespread use in the past years, is unable of representing large and complex domains. That is, the number of random variables associated to the objects in the domain needs to defined in advance. The same is true for the relationships holding between the variables. Therefore, a large and dynamic domain with a varying number of entities and complex relationships evolving in time cannot be properly be represented by a Bayesian network, which are thus propositional in nature. We highlight here the capabilities of the recent Statistical Relational Learning framework of Markov Logic Networks for SA in maritime scenarios. We here briefly provide essential background notions of MLNs, while full details can be found in [11].

An MLN provides an explicit way of encoding knowledge and combining logical and probabilistic reasoning. In First-order Logic, a knowledge base KB of logic formulas is satisfiable only if exists at least one world (truth value of atomic formulas) in which KB is true. A MLN relaxes this hard constraint by associating a probability value to the worlds that do not fully satisfy the KB. Therefore, the fewer formulas a given world violates the more probable it is.

An MLN is a set L of pairs (F_i, w_i) where F_i is a FOL formula and w_i its corresponding real-valued weight. The set of all formulas F_i in L constitutes the KB while the weight w_i associated to each formula reflects how strongly the constraint imposed by the formula is to be respected. The weights influence directly the probability assignment: worlds (that is, truth-value assignments of atomic formulas) which satisfy a high weight formula are going to be much more probable than those that do not.

A Markov Logic Network L together with a finite set of constants C defines a Markov network $M_{L,C}$ that models the joint distribution of the set of random (binary) variables $X = (X_1, X_2, \ldots, X_n) \in \mathcal{X}$. Each variable of X is a ground atom (predicate whose arguments contain no variables) and $\mathcal{X}$ is the set of all possible worlds, that is the set of all possible truth value assignments of n binary variables. The network is built as follows:

- $M_{L,C}$ contains one (binary) node for each possible ground atom given L and C; and
- An edge between two nodes indicates that the corresponding ground atoms appear together in at least one grounding of one formula in L. Ground atoms belonging to the same formula are connected to each other thus forming cliques.

A feature f_i is associated for each possible grounding of a formula F_i in L. Each f_i assumes value 1 if the corresponding ground formula is true and 0 otherwise.

The probability distribution over X assuming values x is specified by $M_{L,C}$ is given by:

$$P(X = x) = \frac{1}{Z} \exp\left(\sum_{i=1}^{|L|} w_i \, n_i(x)\right) \tag{1}$$

where $|L|$ indicates the cardinality of L, thus counting the number of formulas of the knowledge base, and $n_i(x)$ is the number of true groundings of F_i in the world x.

$$Z = \sum_{x' \in \mathcal{X}} \exp\left(\sum_{i=1}^{|L|} w_i \, n_i(x')\right) \tag{2}$$

is a normalizing factor often call *partition function*. Given the joint distribution function in Eq. (1), it is possible to calculate the probability that a given formula F_i holds given the Markov Network $M_{L,C}$ as follows:

$$P(F_i | M_{L,C}) = \sum_{x \in \mathcal{X}_{F_i}} P(X = x | M_{L,C}) = \frac{1}{Z} \exp\left(\sum_{x \in \mathcal{X}_{F_i}} w_i \, n_i(x)\right) \tag{3}$$

where $x \in \mathcal{X}_{F_i}$ is the set of worlds where F_i holds.

While Eq. (1) provides the probability of configuration x of truth values for the ground atoms in the Markov Network, Eq. (3) can be used instead to evaluate the probability that a formula F_i (e.g. a predicate representing an event) holds given $M_{L,C}$ where C is composed by observed entities and other constants. This gives a glimpse of the power of the framework, an arbitrary formula that can be grounded can be queried to get the probability of being true. Thus not only the formulas in L but also any logical combination of them that can be grounded in the Markov network can be queried as well. This is extremely important for a SA system where the operator might want to evaluate the truth degree of a new (complex) event or condition as the combination of existing evidence in the KB.

2.4 Hazmat Example Scenario

In this section, the hazmat example scenario and experimental results from [8] are summarized. The scenario describes several cargo ships heading toward a harbour. Some of the ships carry chemical or generic hazardous materials (hazmat), as, for instance, bleach and ammonia, that when combined may cause a severe threat [12]. The ships are assigned berths in a row, and will be in the harbour before others or at the same time.

The entities in our examples are, as shown in Figure 4, *cargo*, *harbour*, *material* and *berth*, which are linked together by the fact that the cargo ship, carrying some hazardous material (*hazMat*, which can be *dangerous* if combined with other sensitive material) is heading (*isHeadingTo*) toward a certain harbour, in which has a berth. The predicate *hasBerth* takes a triplet of harbour, vessel and berth as argument to bound the three classes. The berth has a predicate *adjBerth*, which is important to indicate that two vessels are moored in adjacent berths, and thus are *neighbours*.

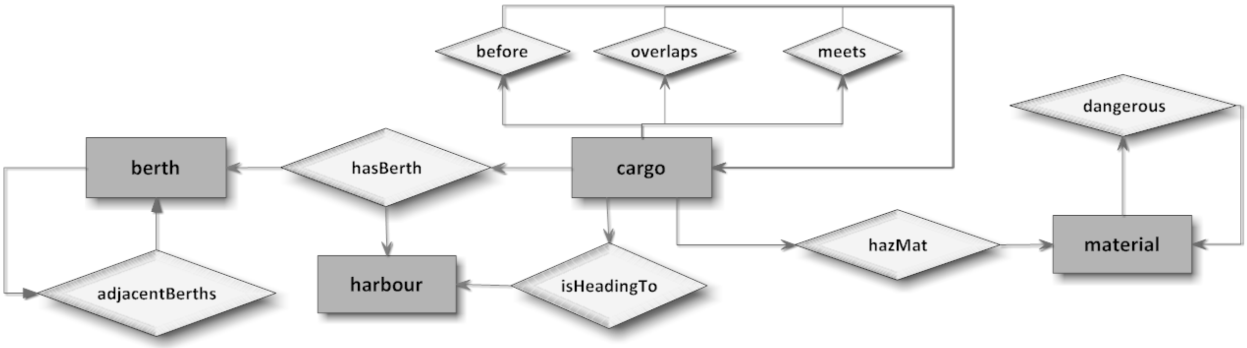

Figure 4: Entities and Relations of the Proposed "Hazmat" Maritime Example [8].

Instead of the seven original predicates (*before*(v_1,v_2), *meets*(v_1,v_2), *overlaps*(v_1,v_2), *starts*(v_1,v_2), *during*(v_1,v_2), *finishes*(v_1,v_2) and *isEqualTo*(v_1,v_2), which define time of permanence at berths of the two ships v_1, and v_2, we shorten the list to *before*(v_1,v_2), *meets*(v_1,v_2) and *overlaps*(v_1,v_2), as these are the most frequent time relations between ships permanence times. In fact, a ship can leave a harbour *before* another comes in, thus the two vessels do not meet. Alternatively, it can stay moored for a long time, which *overlaps* with other vessels permanence. One more case is represented by the *meeting* event, that happens if a vessel leaves just after another one arrives; this situation is relevant as the cargo content may not be fully processed, and still placed on the berth, thus allowing interactions with other ships contents. Other temporal definitions in our domain can be considered special cases of the *overlap* relation. These predicates, that are binary relations, are important as they allow us to properly model the scenario time line and the causality between successive events.

2.4.1 Knowledge Base

The domain knowledge can be formalized with FOL formulas, described in Table 2, where the higher the weight the more confident the statement. The strength of the rules is expressed proportionally to a base weight ω. Weights can be expressed according to experts' knowledge or learned directly from data [11], [8].

Table 2: Knowledge Base for the Hazmat Scenario in FOL with Associated Weights [8].

#	Rule	Weight
1	$overlaps(v,y) \Leftrightarrow overlaps(y,v)$	ω
2	$meets(v,y) \Leftrightarrow meets(y,v)$	ω
3	$neighbours(v,y) \Leftrightarrow neighbours(y,v)$	ω
4	$concurrent(v,y) \Leftrightarrow concurrent(y,v)$	ω
5	$dangerous(m1,m2) \Leftrightarrow dangerous(m2,m1)$	ω
6	$alarm(v,y) \Leftrightarrow alarm(y,v)$	ω
7	$meets(v,y) \vee overlaps(v,y) \Leftrightarrow concurrent(v,y)$	ω
8	$\neg meets(v,y) \wedge \neg overlaps(v,y) \Leftrightarrow \neg concurrent(v,y)$	$4/5\ \omega$
9	$before(v,y) \Rightarrow \neg concurrent(v,y)$	ω
10	$\neg concurrent(v,y) \Rightarrow \neg alarm(v,y)$	ω
11	$cargo(v) \wedge isHeadingTo(v,h) \wedge harbour(h) \Leftrightarrow hasBerth(v,x,h) \wedge berth(x)$	ω
12	$cargo(v) \wedge cargo(y) \wedge hasBerth(v,x,h) \wedge hasBerth(y,z,h) \wedge adjBerth(x,z) \Leftrightarrow neighbours(v,y)$	ω
13	$\neg neighbours(v,y) \Rightarrow \neg alarm(v,y)$	$4/5\ \omega$
14	$cargo(v) \wedge cargo(y) \wedge hazMat(v,m1) \wedge hazMat(y,m2) \wedge \neg dangerous(m1,m2) \Rightarrow \neg alarm(v,y)$	$3/5\omega$
15	$cargo(v) \wedge cargo(y) \wedge hazMat(v,m1) \wedge hazMat(y,m2) \wedge neighbours(v,y) \wedge dangerous(m1,m2) \wedge concurrent(v,y) \Rightarrow alarm(v,y)$	ω

The first six rules (#1-#6) codify the symmetry among elements, and are useful to avoid sorting items. Rule #7 states that two vessels that *meet* or *overlap* are *concurrent* in time, simplifying the concept of "simultaneous" or "operative/moored at the same time". The opposite condition (#8) or the case when one vessel arrives or leaves *before* (#9) others define having no interaction with other ships in the scenario (being not concurrent). If two vessels are not concurrent, they do not represent a threat (#10).

Referring to spatial relationship, a cargo that is heading toward a harbour will have a berth assigned (#11), and two vessels in the same harbour will be *neighbours* only if they will share adjacent berths (#12). If two vessels are not neighbours, they cannot generate an alarm (#13), as well as if they transport cargo materials that are not dangerous when combined (#14).

The main threat can be defined by the rule for which two neighbour cargo ships carry hazmats that are potentially dangerous if combined (#15). In this case, the cargo ships share adjacent berths and are moored in the harbour at the same time.

2.4.2 Contextual Information

Probabilistic knowledge must be integrated with explicit contextual knowledge, as sensory data may be not enough to represent and identify complex situations. A simple low-level anomaly detector would not detect the aforementioned threat, as two cargo ships which enter in a harbour, even carrying hazmats, for commercial reasons raise no alarm. However, additional information provided by context can help to identify a suspicious event.

Context, as described in detail in Table 3, is comprised by scenario-dependent information, which is:

- A harbour H_1 has four berths $B_1,\dots,B_4$, and some of the berths are adjacent. The exact map of adjacent berths can be provided by a human operator. In our case, we codify the proximity with a set of symmetric rules. We suppose, as shown in Figure 5, that berths B_1 and B_2 are adjacent, as well as B_3 and B_4, and B_4 and B_5; and
- Some materials defined by M, if combined together, are dangerous or potentially lethal. This information must necessarily be provided by a SME, as it cannot directly be inferred from materials only. In our example, we suppose that (M_1, M_2), (M_2, M_3) and (M_2, M_4) are dangerous combinations.

Table 3: Contextual Information Provided *a priori* for the "Hazmat" Scenario. Apart from harbour and its facilities, the description of dangerous combinations of materials is provided.

$harbour(H_1)$
$berth(B_1,H_1)$
$berth(B_2,H_1)$
$berth(B_3,H_1)$
$berth(B_4,H_1)$
$adjBerth(B_1,B_2)$
$adjBerth(B_2,B_3)$
$adjBerth(B_3,B_4)$
$\neg adjBerth(B_1,B_3)$
$\neg adjBerth(B_1,B_4)$
$\neg adjBerth(B_2,B_4)$
$dangerous(M_1,M_2)$
$dangerous(M_2,M_3)$
$dangerous(M_2,M_4)$
$\neg dangerous(M_1,M_4)$
$\neg dangerous(M_3,M_4)$

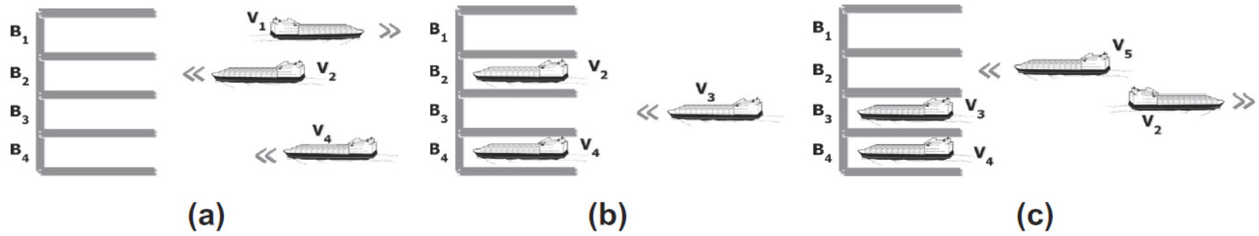

Figure 5: Illustration of the Evolution of the Hazmat Scenario. Cargo ship V_1 leaves much earlier than the arrival of V_2 and V_4 (a), and before V_3 reaches its berth (b). As V_2 leaves, V_5 arrives (c).

As we will see in the experiments, it is important that this information be the most complete as possible, to depict accurately the scenario with its entities and relationships.

2.4.3 Results

We aim to demonstrate how contextual information is a crucial key element to build an exhaustive and accurate situational picture, which allows to timely detect an anomaly.

We imagine a situation as the one described in Figure 5 and Table 4. V_1 leaves the harbour prior to the arrival of V_2 and V_4. After a while, V_3 reaches berth B_3, and it remains there when V_5 arrives and moors at B_2.

Table 4: Observed Facts (Evidence) in the "Hazmat" Scenario. The time of permanence of a cargo at its berth is calculated only with respect to neighbour cargo ships [8].

$cargo(V_1)$	$hasBerth(V_1, B_1, H_1)$
$cargo(V_2)$	$hasBerth(V_2, B_2, H_1)$
$cargo(V_3)$	$hasBerth(V_3, B_3, H_1)$
$cargo(V_4)$	$hasBerth(V_4, B_4, H_1)$
$cargo(V_5)$	$hasBerth(V_5, B_2, H_1)$
$hazMat(V_1, M_1)$	$before(V_1, V_2)$
$hazMat(V_2, M_2)$	$before(V_1, V_5)$
$hazMat(V_3, M_3)$	$overlaps(V_2, V_3)$
$hazMat(V_4, M_4)$	$overlaps(V_2, V_4)$
$hazMat(V_5, M_2)$	$overlaps(V_4, V_3)$
$isHeadingTo(V_1, H_1)$	$overlaps(V_3, V_5)$
$isHeadingTo(V_2, H_1)$	$overlaps(V_4, V_5)$
$isHeadingTo(V_3, H_1)$	
$isHeadingTo(V_4, H_1)$	
$isHeadingTo(V_5, H_1)$	

The fact that a ship is carrying hazardous material and the type of material can be classified as sensory data, as this information can be fetched on-the-fly when the ship becomes a vessel-of-interest or when the system registers the vessel. Also the time predicates can be calculated at runtime, comparing the ETA (Estimated Time of Arrival) and a minimum time of permanence to handle the ship content.

All the cargo ships in our scenario transport hazardous material, but from contextual information Table 3 we know that the dangerous combinations are constituted by (M_1, M_2), (M_2, M_3) and (M_2, M_4).

The query $P(alarm(Vn, Vm)|M_{\{L,C\}})$ represents the probability for predicate *alarm* to be true for a given vessel couple *(Vn, Vm)*, where $M_{\{L,C\}}$ is the Markov Network created by grounding the set of formulas L shown in Table 2, and contextual and sensory evidences are provided according to Table 3 and Table 4 respectively.

Table 5 shows the possible risky combinations of hazardous materials carried by cargo ships that share adjacent berths and are moored in the harbour at the same time. Threats are marked with "Y", while a normal situation is marked with "N" and should raise no alarm. Diagonal terms give no anomaly.

Table 5: Dangerous Combinations of Hazardous Materials Carried by Cargo Ships that Share Adjacent Berths are Boldface and Marked with "Y" [8].

	V_1	V_2	V_3	V_4	V_5
V_1		N	N	N	N
V_2	N		**Y**	N	N
V_3	N	**Y**		N	**Y**
V_4	N	N	N		N
V_5	N	N	**Y**	N	

Hazardous material M_1 is considered dangerous when combined with others, but as the cargo which carries it leaves before others, no alarm is raised. Materials that are brought at not adjacent berths do not constitute a dangerous combination, thus the couple *(V_4, V_5)* does not constitute a threat.

Table 6 and Table 7 show the results obtained for this scenario. In both cases the evidence set is the same, but the contextual information is completely missing in the first experiment. When contextual information is provided, the reasoner sets an alarm in the case of *(V_2, V_3)* and *(V_3, V_5)*, thus matching the truth (Table 5). Contrarily, no alarm is risen when context is missing, as the values for suspicious cargo ships are low.

Table 6: Results for the "Hazmat" Scenario Without Contextual Information [8].

	V_1	V_2	V_3	V_4	V_5
V_1		0.01	0.23	0.17	0.01
V_2	0.01		**0.33**	0.37	0
V_3	0.23	**0.33**		0.34	**0.32**
V_4	0.17	0.37	0.34		0.32
V_5	0.01	0	**0.32**	0.32	

Table 7: Results for Alarm Raising in the "Hazmat" Scenario [8].

	V_1	V_2	V_3	V_4	V_5
V_1		0.05	0	0	0.01
V_2	0.05		**0.95**	0.18	0.01
V_3	0	**0.95**		0.02	**0.89**
V_4	0	0.18	0.02		0.51
V_5	0.01	0.01	**0.89**	0.51	

3.0 IS CONTEXT ALWAYS A GOOD THING? AN INTELLIGENCE CASE

As we have seen in the previous sections, context plays a fundamental positive role in the tasks of achieving results from inference processes about the items of interest. However, at the same time, de-contextualization of objects can be a necessary step in some cases, for example to achieve a successful selection of substitute candidates for tools that are unavailable, or whose presence is undesired. The following is a summarized account of our work on the topic with application in military intelligence, more details can be found in [14].

Hostile *intent*, *capability* and *opportunity* are known to be the three components analysts should look for in detecting potential threats [15]. Given the huge amount of uncertain information to look into, formulating credible hypotheses about potential threats has become even more difficult or impossible in the case of asymmetric warfare where the means for carrying out a hostile plan are often of unconventional type, thus defying all knowledge available from military doctrines.

While intents could be hypothesized based on current intelligence information, capability and opportunity might assume the aspect of normal patterns of life as in the case of recent terrorist attacks (where significant disruption was obtained through non-weapon objects possessing explosive characteristics such as fuel tanks). Since opportunity could be guessed once intent and capability are known, we turn here our attention to determining alternative solutions for assessing capability.

The urgent need for developing automated tools for intelligence analysis [15] should also push the development of alternative ways for encoding and exploiting knowledge in order to facilitate inexact and similarity-based matchings of hypothesized patterns in the knowledge-base. In particular, we propose to extend such ontologies with explicit recordings of physical features in order to capture the intrinsic characteristics that can match function-oriented queries. A mechanism of similarity mapping between ontology classes, using feature-based similarity measures, is discussed to drive the research and retrieval of artifacts which are possible substitute for the proper tool matching the sought after capability. Fusion methods and techniques, exploiting contextual data and information, properly suit for such problems which also often involve soft data issues [16].

Our proposal starts from the analysis of the behaviour and of the ontological status of artifacts. Let's consider the following situation, the so called candle problem, a cognitive performance test, presented by Gestalt psychologist Karl Duncker in his thesis on problem-solving tasks, published posthumous in 1945 by the American Psychological Association. Test subjects are given the materials shown at the top of Figure 6 (a candle, a box of thumbtacks, and a box of matches on a table), and asked to fix the candle to the wall so that, once lit, it will not drip wax onto the table below.

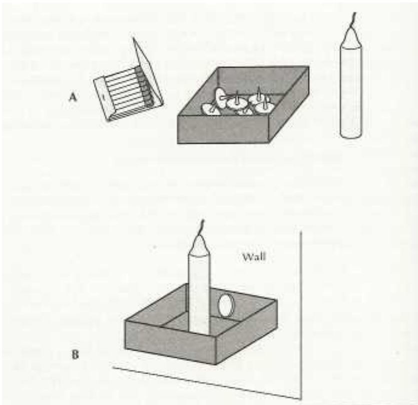

Figure 6: Duncker's Candle Problem.

The test challenges *functional fixedness*, a cognitive bias, which predicts that the participants will only see the box as a device to hold the thumbtacks and generally will not consider it as a functional component independent from the perceived context and therefore available to be used in solving the task. The solution consists in emptying the box of thumbtacks, putting the candle into the box, using the thumbtacks to nail the box (with the candle in it) to the wall, and lighting the candle with the match as in Figure 6 (bottom).

As previously said, the test was created to assess problem-solving skills and the so called "lateral thinking", but we will not deal with its main evaluation purpose, rather we will focus on the mechanism of selection of the functional component, which we define "metaphorical".

3.1 Artifacts and De-Contextualization

The role of context in artifact selection and exploitation is crucial but in a different sense with respect to the role usually played in fusion problems. Context is recognized to be fundamental in achieving tasks by providing expectations, constraints and additional information for inference about the items of interest [2].

On the other hand, in the domain of artifact "metaphors", which involves problem-solving issues, context consolidates functional fixedness obstructing a possible solution as demonstrated by the candle problem.

Moreover, de-contextualization of objects is the first step of a process of "creative" production of substitute tools often deliberately accomplished to perform malicious actions, the most macroscopic among the accomplished ones being the metaphorical substitution "Jet Airplanes" are "Weapons" in the 9/11 Twin Towers attack.

3.2 Capability in Intelligence

An interesting example for the application of metaphorical analysis and reasoning could be the field of military intelligence against asymmetric warfare activities. In general, in the military domain, there is often a more or less well defined "adversary" which could potentially carry on hostile plans. These can be considered a significant *threat* when they meet the threefold condition of: a) being driven by a clear hostile intent, b) being primed by a relevant opportunity, c) being supported by all the capability needed to bring them to completion [15].

Intelligence activities are therefore mainly concerned with assessing adversary intentions with the goal of detecting potential threats as they are being prepared. The task has shown to be particularly difficult in the case of asymmetric warfare where the adversary is purposely not following known military strategies and schemes in order to avoid early detection of own plans.

Given the huge amount of data and information that intelligence analysts have to continuously process from very different sources, there is an urgent need for reasoning methods that can provide automated support to integration and analysis. Shortcomings in the ability to make deductions about missing and conflicting information and the current inability to support automatic context based correlation and reasoning about vast amounts of information are drawbacks to providing a coherent overview of the unfolding events [15].

In the case of asymmetric adversaries, this is complicated by the fact that hostile plans are not only covert but also carried out by unconventional means. This is particularly true in the case of capability, where adversaries often don't exploit the designed and purpose-built tools, but some other tools whose features simply fit their hostile purposes.

Figure 7 shows on the left the three components of a threat as defined above, and on the right the main processing steps that would be required to assess the capability of a hypothesized threat. The process involves metaphorical reasoning in order to detect possible alternative tools for reaching the hypothesized intent *F*. The process is iterative and involves:

1) For each hypothesized intent (purpose) *F*.
2) Abduce *F* -significant properties.
3) Check context for artifact which maximizes capacity (the *F-object*).
4) If NOT present: Extract from KB next possible candidate with sufficient capacity.
5) Loop to 3) until *F-object* substitute is found or termination criterion is reached.

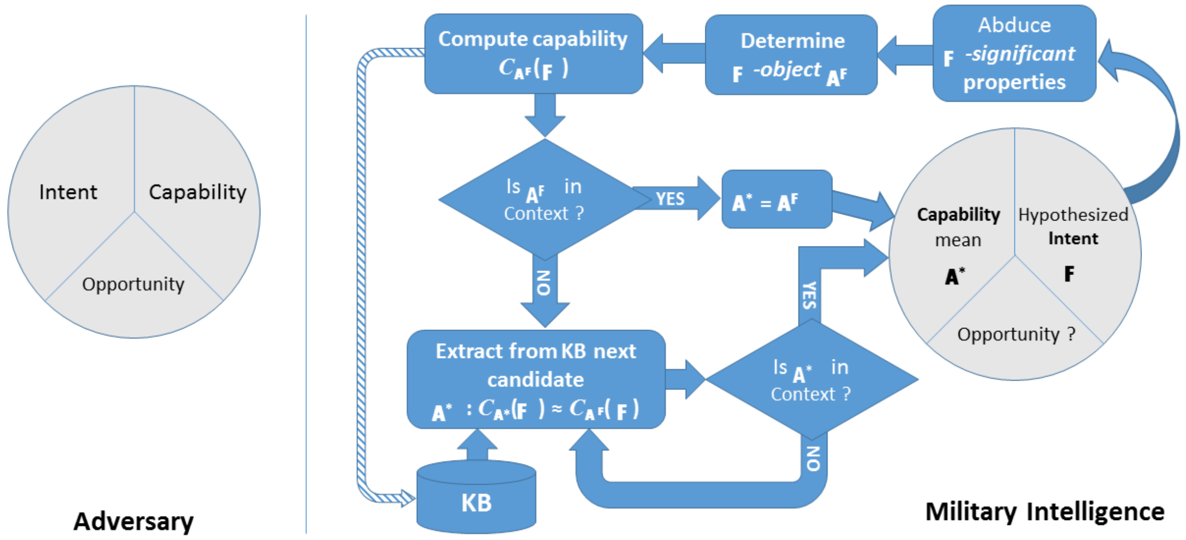

Figure 7: Process of Artifacts' Capability Assessing [14].

The process explicitly looks for the *F-object* that maximizes the capacity, but it could produce a ranking as well and evaluate alternative hypotheses involving tools that have not been explicitly designed for the purpose but that can be used by the adversary as unconventional means.

3.3 Discussion

In everyday life as well as in asymmetric warfare domain, to achieve the intended goals, agents often don't exploit the designed and purpose-built tools but some other tools whose features simply fit for the purpose.

Starting from the ideas carried out to engineer ontologies for functional concepts of artifacts, we propose to extend such ontologies with explicit weighted recordings of physical features. A mechanism of similarity mapping, which will be object of our future research, between instances of property vectors, using feature-based similarity measures, will drive the retrieval of artifacts that are possible candidate substitutes for the proper designed tool.

Context plays a fundamental positive role in the tasks of achieving results from inference processes about the items of interest, namely regarding capabilities related to possible intents but, at the same time, de-contextualization of objects is a necessary step to achieve a successful selection of substitute candidates for tools that are unavailable or whose presence is undesired for example because of a malicious action plan.

4.0 CONCLUSIONS

The exploitation of contextual knowledge has been discussed in IF applications for target tracking, situation assessment and automatic reasoning for military intelligence. While the first two cases provided clear scenarios where the introduction of CI improved system performance by weighting measurements /constraining estimates and refining the inference process respectively, the third one provides an example where de-contextualization allows to overcome functional fixedness and cognitive bias in reasoning.

5.0 REFERENCES

[1] Erik Blasch, Chun Yang, Jesús García, Lauro Snidaro and James Llinas, "Contextual Tracking Approaches in Information Fusion", in L. Snidaro et al. (eds.), Context-Enhanced Information Fusion, Advances in Computer Vision and Pattern Recognition, Ch. 4, pp. 75-99, 2016, doi:10.1007/978-3-319-28971-7_4.

[2] L. Snidaro, J. Garcia, J. Llinas, "Context-based information fusion: a survey and discussion", Information Fusion, Vol.25, September 2015, pp.16-31, doi:10.1016/j.inffus.2015.01.002.

[3] L. Snidaro, I. Visentini, "Context for dynamic and multi-level fusion", in L. Snidaro et al. (eds.), Context-Enhanced Information Fusion, Advances in Computer Vision and Pattern Recognition, Ch. 16, pp. 433-453, 2016, doi:10.1007/978-3-319-28971-7_16.

[4] J. García, L. Snidaro, J. Llinas, "Architectural Aspects for Context Exploitation in Information Fusion", in L. Snidaro et al. (eds.), Context-Enhanced Information Fusion, Advances in Computer Vision and Pattern Recognition, Ch. 8, pp. 187-205, 2016, doi:10.1007/978-3-319-28971-7_8.

[5] A. Steinberg and G. Rogova, "System-Level Use of Contextual Information", in L. Snidaro et al. (eds.), Context-Enhanced Information Fusion, Advances in Computer Vision and Pattern Recognition, Ch. 7, pp. 157-183, 2016, doi: 10.1007/978-3-319-28971-7_7.

[6] I. Visentini and L. Snidaro, "Integration of Contextual Information for Tracking Refinement", *Proceedings of the 14th International Conference on Information Fusion*, July 5-8, 2011, Chicago, IL, USA.

[7] G. Battistello, M. Mertens, M. Ulmke, W. Koch, "Context Exploitation for Target Tracking", in L. Snidaro et al. (eds.), Context-Enhanced Information Fusion, Advances in Computer Vision and Pattern Recognition, Ch. 12, pp. 297-338, 2016, doi: 10.1007/978-3-319-28971-7_12.

[8] L. Snidaro, I. Visentini, and K. Bryan, “Fusing uncertain knowledge and evidence for maritime situational awareness via Markov Logic Networks”, Information Fusion, Vol 21 (January 2015), 159-172.

[9] L.M. Munson, Looking for anomalies all in the wrong places, in: USNI Proceedings Magazines 135/7/1,277.

[10] L. Snidaro and I. Visentini, “Network methods and plan recognition for fusion in crisis management”, in Fusion Methodologies in Crisis Management, Galina Rogova and Peter Scott eds., Springer, pp. 287-312, 2016.

[11] M. Richardson, P. Domingos, Markov logic netwoks, Machine Learning 62 (2006) 107–136.

[12] J. Roy, M. Davenport, Categorization of maritime anomalies for notification and alerting purpose, in: NATO Workshop on Data Fusion and Anomaly Detection for Maritime Situational Awareness, La Spezia, Italy, 2009.

[13] J. Llinas, C.L. Bowman, G.L. Rogova, A.N. Steinberg, E.L. Waltz, F.E. White, Revisiting the JDL data fusion model II, in: Proceedings of the Seventh International Conference on Information Fusion, vol. II, Stockholm, Sweden, 2004, pp. 1218–1230.

[14] G. Ferrin, L. Snidaro, and G.L. Foresti, “Artifact Metaphors: Gaining Capability Using “Wrong” Tools”, *Proceedings of the 17th International Conference on Information Fusion*, July 6-9, 2015, Washington, D.C., USA, pp. 1176-1181.

[15] J. Biermann, P. Horling, and L. Snidaro, “Experiences and challenges in automated support for intelligence in asymmetric operations,” Journal of Advances in Information Fusion, vol. 8, no. 2, pp. 101–118, December 2013.

[16] G. Ferrin, L. Snidaro, S. Canazza, and G. Foresti, “Soft Data Issues in Fusion of Video Surveillance,” in Proceedings of the 11th International Conference on Information Fusion, Cologne, Germany, June 2008, pp. 1882–1889.