Context-Enhanced Information Fusion

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

lauro.snidaro@uniud.it

ABSTRACT

The understanding and principled exploitation of contextual information in information fusion systems is still very limited. Domain knowledge is generally acquired from an expert and applied to stove-piped solutions that can hardly scale or adapt to new conditions. However, research on context-aware fusion system is receiving growing attention, since successful context exploitation can dramatically increase the performances of a modern information fusion systems: from signal processing to intent estimation. In these notes, several context exploitation dynamics and architectural aspects peculiar to the fusion domain are presented and discussed.

1.0 INTRODUCTION

"Context-awareness" might have entered the mainstream language and gained popular attention thanks to the ever increasing adoption of smart portable devices making their way into every activity of our lives. This has happened to the point that the concept of "smart" app (or device) has been paired if not surpassed by the concept of "context-aware" app (or device). The importance of exploiting contextual information (CI) for adaptively react to the surrounding environment is something that has been widely researched in the fields of distributed and mobile computing.

In the past few years, it has become more and more apparent that fusing data from multiple sources is probably not enough. Even combining a potentially large number of sensory sources could lead to some surprises if the value being estimated, the error characteristics of the sources, and the fusion process itself, are not properly referenced to their current relevant context. For example, the state of a target is very likely to have some form of dependence on the environment surrounding the target, on the relationships and interaction with nearby entities, maybe the time of the day and weather conditions. Think for example of the case of estimating position and speed of a car driving through city traffic. The state of the car is going to be influenced by bends and turns of the road, condition of the asphalt, traffic signs, state of preceding and following vehicles, traffic rules, road works, overall traffic conditions, time of the day, weather conditions, and so on.

As already understood in the mobile computing world since at least fifteen years, context is more than just location. However, a clear definition of it is far from being commonly agreed upon. Mostly because of the diversity of the research fields in which the term "context" is picking up popularity. Very loosely speaking, context-awareness involves considering, representing, and exploiting information and knowledge that does not characterize the focal element(s) of interest but the surrounding environment or current situation.

Notwithstanding this growing popularity, the context-awareness concept, namely considering, representing, and exploiting information and knowledge that does not characterize the focal element(s) of interest but the surrounding environment or current situation, had not crossed the borders of the aforementioned computing domain until the past few years. An area that has lately shown a rapidly escalating interest in CI is

Information Fusion (IF). IF systems are traditionally designed to exploit observational data and a priori models and to work well in what can be defined as "well-behaved conditions". However, they cannot be expected to work in problems where the "world-behaviour" is very complex and unpredictable without hard-coded knowledge, or in problems where contextual influences are important or even critical. The development of context-based fusion systems is an opportunity to improve the quality of the fused output and provide domain-adapted solutions.

The understanding and principled exploitation of context in fusion systems is still very limited. Domain knowledge is generally acquired ad hoc from an expert and applied to stove-piped solutions that can hardly scale or adapt to new conditions. However, "context" should play a vital role at any level of a modern fusion system (taking as reference the JDL-Joint Directors of Laboratories- framework [8]): from object recognition through physical context exploitation, to intention estimation through linguistic communication analysis. It would be the key element to gain adaptability and improved performance [1].

The present document aims at providing a concise account of some recent research on context-based IF systems, for a more detailed account the reader is referred to [2], [3].

These notes are structured as follows: Section 2.0 provides a few relevant definitions of context that can be found in the literature; Sections 3.0 and 4.0 discuss how context can improve fusion processes at all levels and how it can help in developing dynamic multi-level solutions; Section 5.0 provides some fundamental concepts and current research topics; Section 6.0 briefly discusses the challenges for building a dynamic and adaptive context-aware IF system.

2.0 DEFINITION OF CONTEXT

While many definitions exist in the literature, intuitively, CI could be said to be information that aids in understanding the (estimated) situation and also aids in reacting to the situation, if a reaction is required. Devlin [4] takes this view, defining context as follows: "a feature *F* is *contextual* for an action *A* if *F* constrains *A*, and may affect the outcome of *A*, but is not a constituent of *A*". Contextual premises can thus be seen as a set of constraints to a reasoning process about a situation. There are of course other definitions of this somewhat slippery term, such as that offered by Dey and Abowd [5], who state that context is "any information (either implicit or explicit) that can be used to characterize the situation of an entity". These definitions imply that these contextual premises are constraints to other premises that could be called "focal" to the formation of our argument or conclusion [6].

A formal investigation on the definition of context is carried out by Rogova and Steinberg in [7], where in addition to describing the concepts of *contex-of* and *concept-for*, challenges are discussed such as utilizing context for understanding of natural language data, context dynamics, context recognition, and contextual reasoning under uncertainty.

Steinberg takes a very significant approach for Situation Assessment (SA) systems in [17], where he considers context as a situation. That is, a set of relationships that can be used either to condition expectations or to improve the understanding of a given inference or planning/control problem. In particular, context for an inference problem is the situation selected by some agent in understanding or solving the problem.

As discussed in [4], the following perspective will be adopted here: in many problems involving interpretation and the development of meaning, there is often some focal data that is purposely collected to help in developing such understanding – in a surveillance application these are the sensor data and possibly human-based observational data. Through analysis, these data can support the formation of what we will call "focal premises" – statements (propositions) about some aspect of the "condition or situation" of interest. To the extent that separate contextual data or information are available, they too can be analysed to form additional premises ("contextual premises") that, fused with the focal premises, can lead to the formation of refined conclusions about the elements being observed. Further definitions of context for IF can be found in [2], [6], [7].

3.0 MULTI-LEVEL MULTI-CUE FUSION

As discussed in [9], modern multi-source multi-sensor fusion deals with different data types that have to be merged together in a coherent decision support process for providing better estimates or more precise operational picture description. It is generally agreed that any process can benefit from a combination of data coming from different and complementary sources, and that this benefit involves at the low-level and the high-level fusion operations.

Multi-modal systems, even more often used for biometric identification and recognition, or multi-sensor multi-cue approaches, which fuse heterogeneous data provided by multiple sensors, are deemed to be a powerful tool to provide a more robust response and enhanced awareness. For instance, in a maritime scenario an AIS contact can be associated with a detection provided by space-based SAR snapshot, and the resulting estimate could be fused with a track obtained by over-the-horizon (OTH) radar. In a surveillance scenario, an IR processed contact can be associated with a detection provided by an RGB camera snapshot, and the resulting estimate could be fused with a track obtained by radar. In these cases, the fusion of multi-sensor information can resolve many data-related issues, as detection imprecisions and difficulties in object association.

Evolving from the multi-sensor concept, the multi-cue approach couches the requirements of representing the object of interest with a collection of distinct and complementary pieces of information at multiple abstraction levels. Fusing information is a matter of reducing the complexity of the domain we want to observe, obtaining an information space of tractable size, and preserving the integrity and the semantics of the data. However, the richness of the domain can be prohibitive to encapsulate in a single piece of information. To give an example, a single sensor can be limited with respect to a space and time, a single feature can be insufficient to describe the appearance of an object, and one sentence to depict a certain situation can be not sufficiently descriptive.

Fusion must consider a multi-cue approach, where the object of interest is described by more than one aspect at once. Moreover, the information we obtain from sensors, wherever they are deployed, is often imprecise, incomplete, ambiguous, unresolved, deceptive, hard-to-be-formalized, or contradictory. Fusing and comparing redundant and heterogeneous data could allow for improvement in addressing the challenges in various environments. In fact, it is deemed that heterogeneity of input data can improve the accuracy and the robustness of an inferencing system.

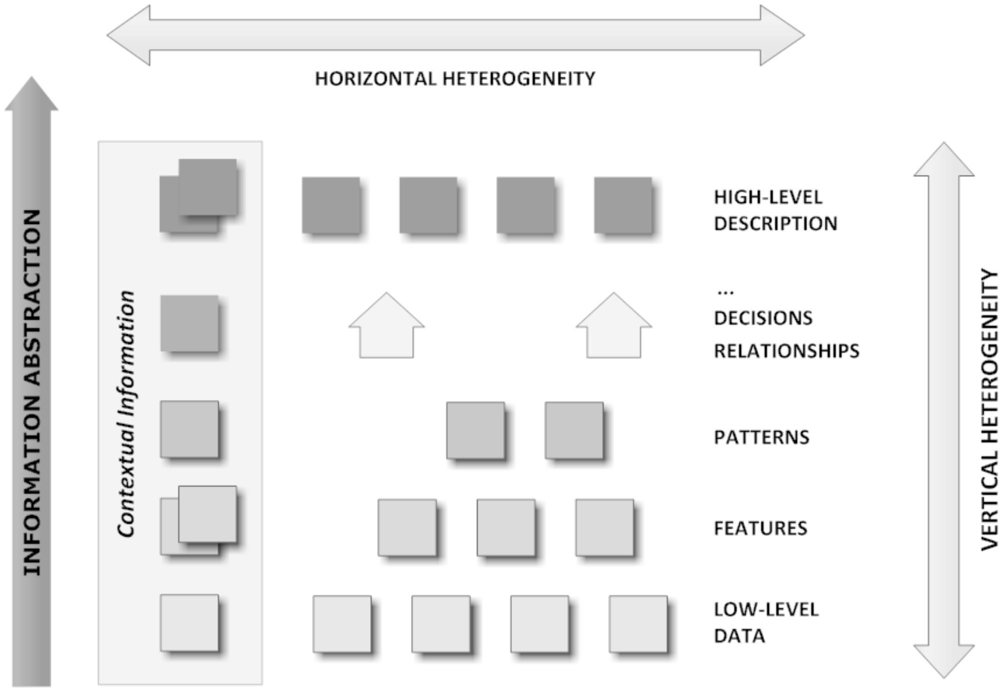

Figure 1: Exemplification of Heterogeneity at Different Abstraction Levels [9].

It is evident that the multi-sensor multi-cue approach forces the fusion process designer to consider non homogeneous data. In particular, we can observe two types of heterogeneity, as shown in Figure 1.

The first is "horizontal", in the sense that refers to same-level information produced or extracted by concurrent or cooperative sources which operate in the same environment. For instance, an object appearance can be described by features as color, shape, velocity or class, which are at the same abstraction level.

The second is "vertical", in the sense that considers that fusion can take place between different levels of information abstraction, from sensory data to features, from features to decisions, including also high-level layers. In this case, the object of interest can be represented by multi-layered information with different degrees of refinement and detail. In a fusion process sense, "complementary" can be taken to mean "associable", and refers to a single given entity, which could be a physical object, an event or a situation. The higher the levels, the more general and abstract the information.

This general taxonomy can be seen as a low-to-high refining process, from a low-level fusion that involves raw data coming straight from sensors to high-level combination of abstract and processed information. The flow through the levels can also be generally considered as a fan-in tree, because each step provides a response to the higher levels receiving the input data from the previous ones. However, it should be noted that sometimes it is possible to merge non-adjacent levels. Relating and associating such data requires some type of semantic transformation or some strategy by which to develop a semantic relation in order to associate them (e.g., ontology).

Frameworks for fusing heterogeneous information have been proposed in the past, but no architecture has become a standard yet, mainly due to the complexity of the problem. According to the previously described taxonomy of complementarity, each level can be considered as a black box, in which different fusion methods can be transparently employed for combining input information.

Fusing intra-layer information is a transparent and stand-alone process: multiple and diverse techniques can be alternated without influencing the other layers input/output. For instance, several classifiers working on different features can be combined by voting techniques, in another situation boosting techniques can be used instead [10].

Inter-layer fusion is also promoted with context: low-level data can be fused with high-level information if properly formatted [11]. For instance, the position of a moving car can be fused with on-line reports or ancillary information on its status to assess if potentially in a threatening situation or compared with a city map to signal the risk of violation of sensitive areas.

To provide a concrete example, we consider a tracking system, which has to cope with unidentified objects detected by physical sensors and with uncertain information coming from human intelligence:

- The data level can be constituted, for instance, of different measurements of the object position provided by different sensors (geo-location given by camera, radar, etc.);
- The feature level comprises different object characteristics, basic pieces of information (Haar features output, color, class, etc.) which describe the object's physical appearance;
- The next step includes the contextualization of the object in a network of entities and relationships, which may be built on a priori knowledge or obtained by data mining. Entities involved in a urban domain, for instance, have complex relationships and are constrained by the surrounding environment; and
- One level above we can find trajectories and historical motion patterns, which are structured pieces of information obtained by exploiting features and data at lower level. From trajectories and from formatted knowledge provided by operators or intelligence, we move one step up, and include information regarding the object's capabilities, its behaviour or even its future intent.

3.1 "Multi-Level" is not "Hard+Soft" Fusion

To further clarify the nature of the data/information that a multi-level fusion is called to process, it can be useful here to further stress the fact that "hard" or "soft" data, that is coming from sensors or humans, are not necessarily processed by low and high level fusion processes respectively. In fact, the level of the data and the type (hard/soft) represent different dimensions and Table 1 exemplifies the four combinations that can hold.

Table 1: Low and High Level Data vs. Hard and Soft Sources.

	Low-Level	High-Level
Hard	Typically raw numerical data such as positional data provided by sensors. However, data can also be text strings as output of a classification algorithm (i.e. labels).	Sensor or system outputs with high semantic value such as detection of events or situations which are typically underpinned by relations holding between the elements involved in the detected pattern (e.g. relations among detected entities, relations between entities and context, etc.).
Soft	Typically numerical information (e.g. number of observed entities, distance, etc.) or text labels regarding entities.	Semantically rich observations typically couched in natural language.

In particular, what is referred to in Table 1 as "Low-level" fusion comprises JDL levels [8] 0-1, while "High-level" is considered level 2 and above [11]. It is true that soft sources typically provide high level information (e.g. reports on behaviours or relationships patterns of observed entities), and that hard sources typically provide low level numerical data or classification labels (e.g. measurements of physical quantities, attribution of a class label to an instance of that class, etc.). However, the opposite is also possible: humans providing rough or fuzzy estimates of quantities (e.g. “about 100 meters”) or a surveillance software proving high-level information such as the detection of complex patters of activity (e.g. “person *p1* and person *p2* entered car *c1* at *11:55*. Car *c1* left the parking lot at *12:00*”[1]). The distinction between hard and soft sources is made principally because of the different models for uncertainty, reliability, accuracy, and trust needed for characterizing the source.

4.0 CONTEXT AS BINDING ELEMENT FOR MULTI-LEVEL FUSION

IF processes and systems have been traditionally designed to fit low-level observational data with a priori models, often at one IF level only. These approaches have shown to work well for static and predictable domains where a priori knowledge is rich, but are not fully capable to solve problems in which the global behaviour is very complex or multifaceted, or where contextual influences are evident.

This is particularly true for surveillance systems as, for instance, maritime situational awareness with an emphasis on counter-piracy, counter-terrorism or harbour/port protection [13]. In this case, CI such as traffic condition, restricted areas, harbour maritime rules, human operator inputs, etc. constitutes essential information to improve several aspects of surveillance.

The issues above can be addressed exploiting contextual knowledge to enhance the performance and capabilities of a fusion system. An integrated approach should extend problem-space modelling with a structured multi-level representation of semantically rich domain knowledge, which describes situations and possible entities interactions of the observed scenario.

From the literature, high-level information is not often merged with sensory data, even if its richness and completeness are extremely useful to properly interpret the available stream of data. Qualitative high-level knowledge can help to infer hidden states from low-level data generated by sensors, other fusion processes or human reports. In other words, it can reduce uncertainties in problems where normally experts would need to be consulted.

The main concepts remarked here are the following [9]:

- Context allows for improving the associability between problem-space knowledge and models and observational data, increasing fusion performance. The a priori model can be better fit to data exploiting the semantics provided by CI.
- Context can be provided at different IF levels. Different types of context, from coarse to refined information, can be injected at different stages of the fusion process, overlapping or overriding each other.
- Context is an important source of semantics, which provides a means to bind data and models.

5.0 CONTEXT IN FUSION: SOME FUNDAMENTALS

In the following, some fundamental concepts on context are highlighted, detailing a few specificities that need to be taken into consideration when developing context-aware fusion systems. In particular, the concepts of internal and external context can be found across many domains and we discuss here their implications in fusion systems for SA [2].

[1] The message is here represented informally for the sake of simplicity. Controlled languages such as BML [11] can instead be used as formal representation and communication means. A human can also provide input to the system in a formal language through a proper human-machine interface.

5.1 External and Internal Context in Cognitive Sciences

Among the different ways in which context has been modelled, the partition between *external* and *internal* context is a concept that appears to be widely accepted [2]. In Ubiquitous Computing and earlier in Cognitive Science, internal and external context are meant to be two different dimensions associated to the external physically measurable world and the internal (unobservable) state of the user including goals, tasks, emotional state, etc. This definition appears to be an adaptation for the domain of pervasive computing and context-aware devices of what was meant by earlier works in the field of cognitive science and perception. In particular, Kokinov [18] states that “context is the set of all entities that influence human (or system’s) behaviour on a particular occasion, i.e. the set of elements that produce context effects”. Then he describes, citing quite a few references in cognitive science dating back to years 1986, 1988 and 1993, the notions of external and internal context where the former refers to “physical and social environment or the setting within which the subject’s behaviour is generated” while the latter “subject’s current mental state within which the subject’s behaviour is generated’”. In this domain, external context is then seen as the sphere of subjective perceptions of the surrounding environment that have an effect on the subject’s mental state. According to these definitions, the *External context → Internal context* relation appears predominant where external factors produce effects on the internal context of the subject. Although the relation is not strictly in that direction only, as the internal context (e.g. mental state) of the subject can influence the correct or complete perception of the surrounding environment and thus in turn its influence can be for example, with different degrees of consciousness, altered or even prevented (Figure 2).

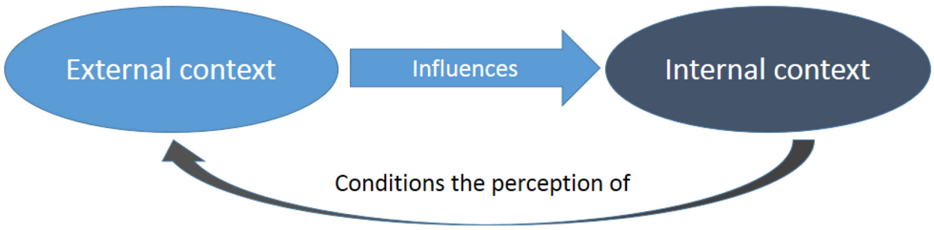

Figure 2: Relations Between External and Internal Context in Cognitive Science [2].

As already mentioned, these notions of internal and external context developed in the cognitive sciences domain have been quickly adopted by the researchers in mobile and pervasive computing where external attributes, most notably location, are sensed in order to provide relevant information to the user. Most of the papers concentrate on the exploitation of external context since some attributes of it can be sensed by low-cost hardware. Even though being generally non-observable, there are still some good chances of guessing the internal context of the user (of course a subset of it). Save for the cases where the user is directly providing (part of) it, for example by sharing her/his emotional state explicitly or by disclosing interests or intentions by searches in search engines, other mechanisms involve for example the analysis of web navigation patterns, opened documents, etc.

5.2 External and Internal Context for Information Fusion

We have seen how the concepts of external/internal originated in the field of cognitive science, and we would like now to highlight what are the characteristics, commonalities and differences in typical tasks in the IF domain.

5.2.1 Uncertainty

Even though almost all of the domains can be seen in terms of IF as soon as multiple sources of data/information are present and there is the need to combine their products in order to obtain better estimates of a certain variables, typical IF systems and applications generally have the common problem of lack of direct information from the focal entities of interest. SA systems, for example, have to go through a number of processing steps, also combining heterogeneous data, in order to estimate the status and intentions (or purpose) of non-cooperative entities (or process/system). In addition, observations from sensors are generally noisy and sources of information can have different level of trust and provide outputs with different quality, therefore making fusion a real necessity. Due to its inherent relational nature, the representation of context uncertainty can follow the models developed for situations and reasoning about situations [13], [17], [7], however an additional level of uncertainty is given by the *relevance* of a context for a given task [17] which is still an under-researched topic.

5.2.2 Observed and Observable Context

Here we consider a working definition of a *situation* as the collection of all the entities, their attributes, the relations among them and the environment, and the events occurring in a given scenario at a certain time. The entire real situation, giving a perfect account of what is happening in the scenario, is of course impossible to observe and represent. A SA system can only observe a subset of the real underlying situation and this subset is given by the purpose of the system, that is the system's internal context. This means that if the system was designed for a specific purpose or if its current setting is directed to a certain objective, then the system is configured for observing a specific subset of the situation. This is in practice complicated even more by the sensing and interpretation capabilities of the system, the noise corrupting the observations, uncertainties involved in the processing algorithms, etc., making the situation actually observed an even smaller subset of what the system was intended to perceive and understand. Since the purpose of the system and its inner workings are known to the system designer, the internal context of the system is here understood as totally observable. These concepts are graphically described in Figure 3.

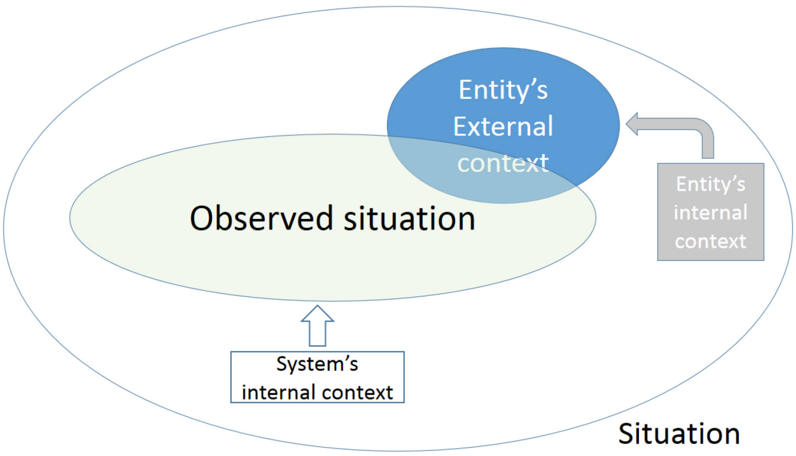

Figure 3: Situation and Context [2].

5.2.3 External and Internal Context

An interesting notion is given in [16] is the following "context and situation can only be defined with respect to an entity for a given purpose". A similar understanding is given by Steinberg in [17], stating "the particular concerns of some agent (e.g. a person or an automated inference system) determine which situations are under consideration as context for those concerns".

The authors in [16] generalize that the context at time t that relates to a set of agents for performing a task is a composition of situations in a time interval between a starting time t_0 and t. A situation is defined as the set of observed values of the variables that relate to a given agent for performing the given task.

The specific characteristics of fusion systems for SA so far described bring us to a revision of the concepts of external and internal context as follows. If a given entity *f* in the scenario is to be considered of interest, that is the focal element, then the external context of *f* is understood here as a subset of the situation that can be put in relation *f* because of its internal context. That is, the current goals and objectives of *f* define what could be considered as contextual for it at a given time as shown in Figure 4, thus in a sense reversing the direction of the main conditioning relation of Figure 2.

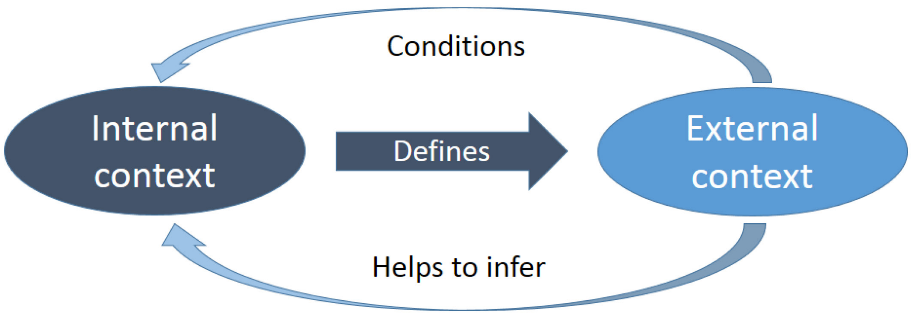

Figure 4: Typical relations between external and internal context in fusion for SA where the mission of the system is typically to infer the goals/purpose (internal context) of a focal entity. The internal context of the focal entity defines what in the current situation is contextually relevant to it (external context). This external context can be observed or inferred by the system in order to discover the entity's internal context. At the same time the external context (e.g. road network) conditions the goals of the focal entity [2].

Here, the goals of *f* (unknown and to be discovered) project a number of relations to elements of the situation that are relevant to *f* 1) for accomplishing its goals or 2) relevant to the system because their contextual effects on *f* help to understand the behaviour of *f* and infer its goals or purpose. This means that in a fusion system for SA both the focal's goals and CI need to be continuously estimated in an iterative process: the initially hypothesised goals of *f* define what is contextual to *f* that in turns helps to refine/confirm/reject the initial hypothesis. This proceeds continuously and dynamically as *f* can change its goals over time, also depending on the focal's own contextual knowledge and perception of it that can influence/change its own internal context (Figure 4).

With respect to [16], the context is not intended here as a composition of situations since the valid external CI is here understood as being the one that relates to the current goals of the focal. CI that has exited the current scope of validity is treated as historical context that can be used for automatic context learning purposes.

5.3 *A Priori* Knowledge, Context, and Adaptability

First of all, a first step that can help the design phase of a context-aware fusion system is a clear separation between what should be considered as an always-valid static a priori knowledge and what could instead be relevant to the specific domain at hand. This allows to develop knowledge repositories in a modular way, thus allowing easier context switching and system re-deployment. Also understanding what kind of knowledge is of dynamic nature allows to develop systems that can automatically exploit a change in the observed environment. This understanding brings in the concept of context sifting described in Section 6.1. Channelling CI to the proper algorithms can easily be done at design time (*a priori* based) or while the system is off-line with the help of a human expert [4]. More complicated is the case of on-line context exploitation (*a posteriori* based) [4]. However, it should be noticed that the sifting procedure would work differently depending on the type of context received by the system while on-line. If context is received as value of a known property/attribute, then the value is exploited by the system according to the existing system design. Take for example the case of weather conditions: this is dynamic contextual knowledge that can be provided both by human observers and sensors. Supposing that weather can take values in a pre-defined set of strings like {Rainy, Stormy, Sunny, Cloudy}, then on-line processing of this CI occurs as per the existing architecture. If context is instead provided, typically by a human expert, as a new relation or rule such as *weather(area,Rainy)^sensorType(s,Camera)^covers(s,area) →confidence(s, Low)*, then this kind of knowledge could simply be added to the contextual knowledge base and immediately produce effects on entailed predicates. A much more advanced architecture would need to reason on CI, system capabilities, and scenario entities in terms of causes and effects as mentioned in Section 5.5.

5.4 Context Heterogeneity and Information Fusion Levels

CI can be represented with different and inhomogeneous information tokens; each of these depicts a snippet of background, which is associated to a situation or an event, or a physical object. Heterogeneity is a fundamental requirement in fusion, since it provides complementary, redundant, rich information at different levels of abstraction, which can be helpful toward understanding the focal states of interest or for describing these states, or other purposes. No single sensor is currently capable to capture the complexity and vastness of a domain, but expert knowledge and, more in general, high-level soft (e.g. text) data should be blended together with low-level hard (e.g. physical measurements) information to obtain a more comprehensive, timely and accurate situational awareness picture.

A synergistic exploitation of available contextual pieces is demanded, both for intra-level and inter-level fusion: fusion can take place between information produced or extracted by concurrent or cooperative sources which operate in the same environment, and between different abstraction levels as well [9].

5.5 Middleware

A context-aware fusion architecture needs to be designed to be able to incorporate the capabilities above mentioned to properly exploit contextual knowledge both for understanding better the observed domain entities and to automatically regulate the system itself to cope better with changes in the environment. A form of middleware layer would be required as the means of access to the contextual databases and of binding system capabilities to context as well. This layer would also be responsible to perform the context sifting and switching capabilities described in this chapter. However, a significant additional effort is required to properly formalize such an architecture and the representation of all aspects of CI, some of which have been highlighted in this work. Further insights to the middleware concept for context-aware IF can be found in [19].

6.0 DESIGNING FOR ADAPTIVITY

As we have seen in the previous sections, context plays a vital role at any IF processing level and it allows the possibility for a system to change and adapt to better react to unpredictable or potentially harmful events. The system should not only adapt to changes in observation data or availability of resources, but also to the presence of new or updated CI, such as time of the day or adapting the exploitation of the information a sensor has about the object being observed. This form of adaptation is pertinent to JDL model Level 4 optimization schemes and policies [8]. In the following, concepts relevant to a context-aware fusion system able to adapt to changing context conditions are discussed [9].

6.1 Context Sifting

Contextual elements can be defined in very different ways in a fusion system. As already mentioned in Section 3.0, knowledge about a particular domain can have effects at different fusion levels. It could be of interest to the system designer or from an algorithmic point of view to understand how CI can be decomposed, possibly translated, and exploited at different levels. For example, a surveillance application for the security of an enterprise or public building could have information, provided by a domain expert or automatically acquired, regarding the activity of the parking lot attached. Information could be in the form of overall working hours of the employees in the building. This could suggest the typical arrival and departure hours for vehicles in the parking lot thus providing a way for detecting unusual activity outside of the normal working hours time window. This kind of information would therefore be exploited at high level (JDL Levels 2 and 3) in a fusion system.

Other information could comprise the map of the area, entrance and exit gates, allowed parking slots, etc. This kind of information could be useful to constrain vehicles' movements in that area and thus be exploited by Level 1 algorithms for target tracking. At the same time, this information could be exploited at high level to detect unusual movement patterns. The simple parking-lot example shows how contextual knowledge could affect a different number of levels depending on the entities involved. The process by which the entirety of the domain knowledge acquired is fractured into pieces and assigned to the proper algorithms is here referred to as *context sifting* and illustrated in Figure 5.

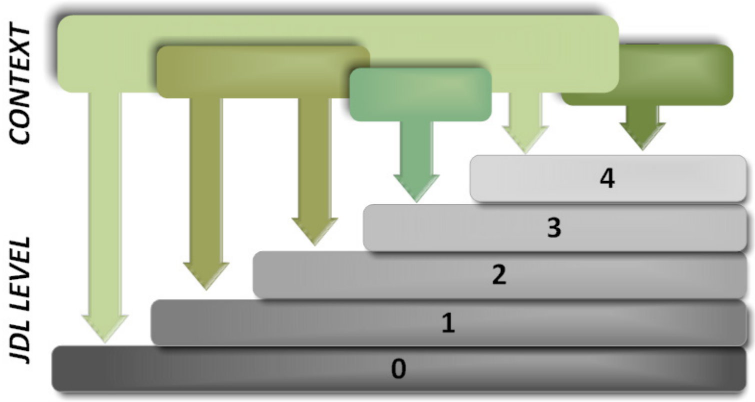

Figure 5: Sifting of CI to Processes at Different JDL Levels [9].

The problem is relevant for both off-line repositories and on-line flows of information that could carry relevant cues of the domain at hand. While in the former case contextual factors could be assigned by the system designer a priori, the mechanism for sifting CI a posteriori is particularly interesting [4]. In both cases, CI can be provided by collected sensory data (hard data) or expert knowledge (soft data).

Figure 6 illustrates a possible process by which domain knowledge, possibly in the form of soft data reports, is first stored in computable form (knowledge structuring) and then actually sifted to the proper algorithms.

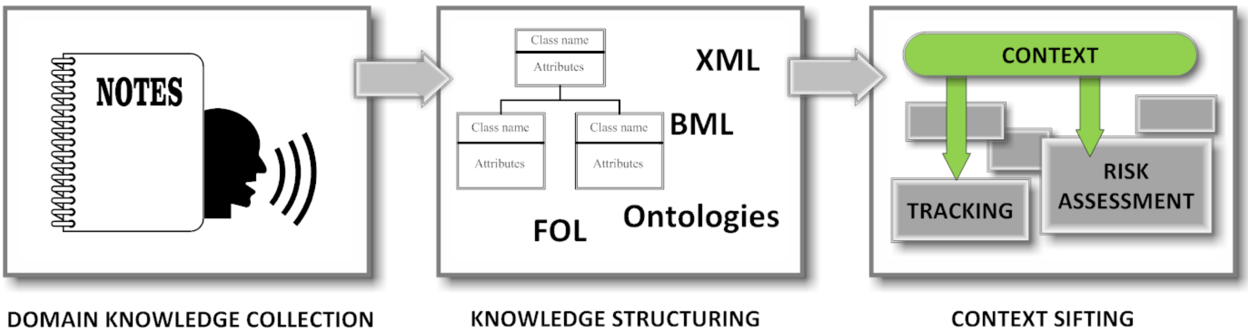

Figure 6: Representation and Sifting of CI [9].

Rich domain knowledge can comprise descriptions of entities and their features and relations, rules describing typical entity behaviour, regulations, constraints, uncertainties, etc. Different representations can be used for such wealth of knowledge. For example, in [12] different mechanisms are employed to encode knowledge about a urban environment and road network to road to constrain tracking of vehicles. In [13] the Markov Logic Networks framework was used leveraging both the expressive power of first order logic and the probabilistic uncertainty management of Markov networks to encode high-level knowledge in maritime domain.

The process by which on-line CI is dynamically exploited at different fusion levels has not seen, to our knowledge, practical implementation yet. To dynamically sift incoming contextual knowledge a context-aware architecture would be needed where the internal parameters of the algorithms (e.g. tracking, event detection, etc.) would be bound to concepts expressed in the contextual knowledge repository. A key element for a possible solution could be the development of a middleware processing layer as briefly mentioned in Section 5.5 and more deeply discussed in [19].

6.2 Context Switching

CI, as the name implies, has local scope and validity and is thus pertinent to the scenario at hand. The granularity of the scope of certain information can be more fine-grained and be applicable only to sub-areas of the observed environment. Figure 7 illustrates the case of two cameras observing the activities of two different parking lots.

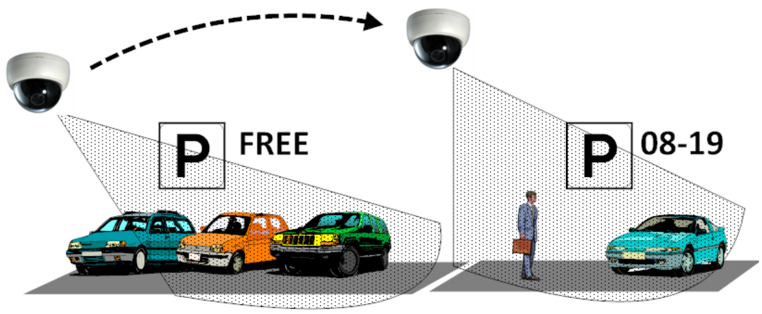

Figure 7: Context Switching, *Shallow Transfer*. Source and target domain are the same. Source and target tasks are the same [9].

Take for example the case of two different parking areas for a public/enterprise building. The first is for visitors/customers and allows free parking while the second is dedicated to staff personnel only. Activity in the two areas is likely to be different, for example knowing that working hours should be in the 0800-1900 time range, the system could exploit this information to detect anomalous events (Section 6.1). This knowledge does not necessarily apply to the first parking area. How much of the knowledge that applies to one area applies also to second one? Given the effort required to build an effective knowledge base (made of expert provided notions or learnt online by the system), from a knowledge engineering perspective it would be interesting to know how much of a certain context knowledge can be reused and applied in another situation. This of course is not only related to the effort or cost in building the knowledge base, it directly determines the capability of a system to adapt to changing domain conditions, and it affects its capability to be redeployed in another scenario.

The context switching capability discussed here is inspired by available literature on *transfer learning* [14]. Transfer learning deals with transferring classification knowledge and capability from a source domain to a target domain or from a source task to a target task. The idea here is brought even further by postulating a transfer of the entire system capabilities from a scene to another or from a domain to another. In Figure 7, the source and target domain are the same (parking lot) and the source and target task are also the same (surveillance). Nonetheless, the two parking areas have different maps, different patterns of activity, different entrance and exit gates, etc. However, in both cases the expected entities are vehicles and pedestrians with their known typical (for such scenario) movement speeds. Also models of usual and unusual behaviour are likely to be transferable between the two scenes. In the transfer learning literature this type of transfer is often referred to as *shallow transfer* [14].

A more drastic context switch is illustrated in Figure 8. Even though source and target tasks are the same (detecting unusual activity) here very few contextual knowledge is likely to be transferred between domains. This is often called *deep transfer*.

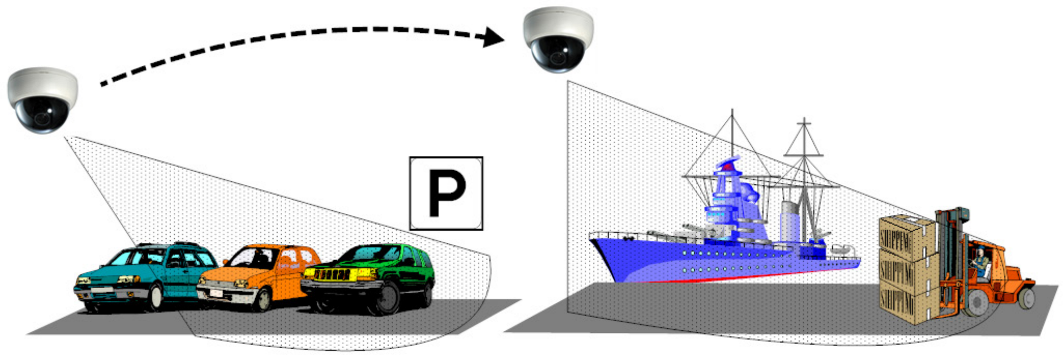

Figure 8: Context Switching, Deep Transfer. Source and target domain are different. Source and target tasks are the same [9].

However, a fusion system with context adaptable architecture should be able to be easily redeployed in a different scenario without much re-training and re-configuring. For example, normal patterns of activity in a surveillance scenario can be re-learnt online. Domain knowledge and rules should be easily be swapped by changing knowledge repositories [13]. Deep transfer can also be performed automatically by inductively learning knowledge from data in a target domain starting from the knowledge base of a source domain. Deep Transfer via Markov logic [15], for example, starts by lifting knowledge bases expressed within the Markov Logic Networks framework to second-order clauses where domain-specific predicates have been replaced by predicate variables. This second-order set of clauses represents domain-independent knowledge that can be exploited as declarative bias, that is a starting point, while learning knowledge in the target domain. This allows to both reduce the computational burden of automatically learning new knowledge and producing a more refined and accurate knowledge base [15].

7.0 CONCLUSIONS

The exploitation of context information is receiving growing interest in the IF community, with a number of works presenting performance improvements via context exploitation in the underlying models.

There is a need for the fusion community to develop an engineering and design methodology for the inclusion and exploitation of CI. In these notes, several key concepts and top-level issues for the incorporation of CI in fusion systems have been briefly presented along with references to recent works.

Research on context for IF can contribute to the development of a new generation of fusion systems, based on solid and general theoretical foundations for a multi-level and synergic interplay between observational and contextual data, beyond the abundant particular cases in current literature.

8.0 REFERENCES

[1] L. Snidaro, J. Garcia, J. M. Corchado, "Context-based Information Fusion", Information Fusion, Guest Editorial, Vol. 21, January 2015, pp.82-84, doi:10.1016/j.inffus.2014.02.001.

[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, J. García, J. Llinas and E. Blasch (eds.), Context-Enhanced Information Fusion - Boosting Real-World Performance with Domain Knowledge, Advances in Computer Vision and Pattern Recognition series, , Springer International Publishing, 2016, doi: 10.1007/978-3-319-28971-7.

[4] K. Devlin, "Confronting context effects in intelligence analysis: How can mathematics help?", Center for the Study of Language and Information, Stanford University, 2005 (unpublished).

[5] A. Dey, G. Abowd, "Towards a better understanding of context and context-awareness", Proceedings of the Workshop on the What, Who, Where, When and How of Context-Awareness. The Hague, The Netherlands, 2000.

[6] J. Llinas, L. Snidaro, J. García, and E. Blasch, "Context and fusion: Definitions, terminology", in L. Snidaro et al. (eds.), Context-Enhanced Information Fusion, Advances in Computer Vision and Pattern Recognition, Ch. 1, pp. 3-23, 2016, doi:10.1007/978-3-319-28971-7_1.

[7] G. Rogova and A. N. Steinberg, "Formalization of "Context" for Information Fusion", in L. Snidaro et al. (eds.), Context-Enhanced Information Fusion, Advances in Computer Vision and Pattern Recognition, Ch. 2, pp. 27-43, 2016, doi:10.1007/978-3-319-28971-7_2.

[8] 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.

[9] 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.

[10] I. Visentini, L. Snidaro, and G.L. Foresti, "Diversity-aware classifier ensemble selection via f-score", Information Fusion, Vol. 28, March 2016, pp. 24–43, doi:10.1016/j.inffus.2015.07.003.

[11] J. Biermann, J. García, K. Krenc, V. Nimier, K. Rein, and L. Snidaro, "Multi-level Fusion of Hard and Soft Information for Intelligence", in L. Snidaro et al. (eds.), Context-Enhanced Information Fusion, Advances in Computer Vision and Pattern Recognition, Ch. 17, pp. 455-479, 2016, doi:10.1007/978-3-319-28971-7_17.

[12] 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.

[13] 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, pp.159-172, doi:10.1016/j.inffus.2013.03.004.

[14] S.J. Pan, Q. Yang, "A survey on transfer learning", IEEE Transactions on Knowledge and Data Engineering 22(10), 1345–1359 (2010), doi:10.1109/TKDE.2009.191.

[15] J. Davis, P. Domingos, "Deep transfer via second-order markov logic", Proceedings of the 26th annual international conference on machine learning, pp. 217–224, 2009.

[16] J. Coutaz, G. Rey, Foundations for a theory of contextors, in: Computer-Aided Design of User Interfaces III, Springer, 2002, pp. 13-33.

[17] A. Steinberg, "Situations and contexts", ISIF Perspectives on Information Fusion, Vol.1 (June 2016), pp.16-24.

[18] B. Kokinov, "A dynamic theory of implicit context", in Proceedings of the 2nd European Conference on Cognitive Science, Manchester, UK, 1997.

[19] 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.