Recent Trends in Context Exploitation for Information Fusion and AI

Lauro Snidaro, Jesus Garcia, James Llinas, Erik Blasch

■ *AI is related to information fusion (IF). Many methods in AI that use perception and reasoning align to the functionalities of high-level IF (HLIF) operations that estimate situational and impact states. To achieve HLIF sensor, user, and mission management operations, AI elements of planning, control, and knowledge representation are needed. Both AI reasoning and IF inferencing and estimation exploit context as a basis for achieving deeper levels of understanding of complex world conditions. Open challenges for AI researchers include achieving concept generalization, response adaptation, and situation assessment. This article presents a brief survey of recent and current research on the exploitation of context in IF and discusses the interplay and similarities between IF, context exploitation, and AI. In addition, it highlights the role that contextual information can provide in the next generation of adaptive intelligent systems based on explainable AI. The article describes terminology, addresses notional processing concepts, and lists references for readers to follow up and explore ideas offered herein.*

Artificial intelligence (AI) was popularized in the 1950s and has since gone through three major instantiations. Typically, AI is meant to emulate human reasoning and planning. The first AI methods trained on simple use cases employing handcrafted knowledge. The second phase of AI focused on computer methods but was limited by insufficient data for training and development. The third phase is based on statistical-based deep learning, which requires many training exemplars and for which, in an era of big data, there is great hope for the realization of advanced capabilities. In all three cases, context extends use cases in the path toward generalization by determining which data were needed and bounding the type of data assessed.

 ISSN 0738-4602

Development of explainable intelligent systems has been identified as a key area of research and a possible major step in AI. The recent DARPA explainable artificial intelligence program is the subject of significant funding and is expected to close by 2021.[1] Presented at the 2017 International Joint Conference on Artificial Intelligence workshop on explainable AI, a survey by Biran and Cotton (2017) indicated AI paradigms whose output can be understood by humans, by introspection or by building explanations. Most current machine- learning techniques are difficult to explain because their models are complex, usually of black-box type, and therefore not easily interpretable. Other classic methods are instead inherently interpretable, as is the case of rule-based systems, decision trees, causal networks, or hidden Markov models. Context-aware systems are mentioned as advanced methods for sensing the environment to enable self-adaptation. With sensing and adaptation, the interaction with users gains in understanding when high-level concepts are managed by the system to explain the changes learned.

This article highlights three concepts that themselves are complex and multidisciplinary: Context information (CI), information fusion (IF), and situation awareness (SAW), each with long histories of research and publications (Snidaro et al. 2016; Snidaro, Garcia, and Llinas 2015). Adding to these foundations, elements of AI related research are organized in relation to the concepts. Context allows systems to augment the observations and enhance meaning, with the goal of developing solutions for inferences contextualized in a domain with mechanisms for adaptation to situations.

In this article we will first present the IF terminology and problems addressed, followed by an introduction of the role of context in IF. Next, we will detail context adaptation from the viewpoint of AI research, highlighting the expectations and open challenges for AI research. We will then present the role of AI in IF problems, with an emphasis on the processes needed for situation understanding, give some perspectives on the current research on AI and IF systems leveraging context adaptation, and finally present our conclusions.

An Introduction to IF

IF is an area of research and development now maturing in certain ways, as well as an area of expanding study. In its early development (in the 1980s), and perhaps typical for emerging areas of science and technology, normalizing the language of the field was an initial complexity, and the status of that language remains somewhat inconsistent and ambiguous. Therefore, this article does not focus on debatable terminology but points out agreed upon concepts. For example, the boundaries between sensor fusion, data fusion, and IF (as well as knowledge fusion) are generalized to IF to provide readers with normalized concepts. The choice is partly motivated by the focus of this article, that is, to discern the concepts, techniques, and applications of IF that use contextual material, because we argue that the categories of information falling into the class of contextual are generally broader and more informative than those of data, especially for sensor data that generally comprise discrete quantitative measurements.

Definition of Data (Information) Fusion

The initial data fusion lexicon, produced by the US Joint Directors of Laboratories (JDL) Data Fusion Subgroup in 1987, defined data fusion as

> a process dealing with the association, correlation, and combination of data and information from single and multiple sources to achieve refined position and identity estimates, and complete and timely assessments of situations and threats, and their significance. The process is characterized by continuous refinements of its estimates and assessments, and the evaluation of the need for additional sources, or modification of the process itself, to achieve improved results. (White 1987)

A simpler version might be "a process of combining data or information to develop improved estimates or predictions of entity states." The point is, data fusion — henceforth IF in the paper — is an automated information process that combines data, in the broadest sense, to estimate or predict the state of some aspect of a problem space of interest and thus to improve those values beyond what could be done without such integration. The notion of estimation should also be clearly understood. Few would argue that the inputs into an IF process are random variables, where sensor responses are typically modeled using statistical concepts to represent the imperfections in sensing operations and the resultant observables or measurements. In the broadest sense the IF process, whatever it does, can be conceptualized as a function, and clearly any function of a random variable yields a random variable; in other words, the fused estimate, no matter how elegantly calculated, is a random variable having a statistical distribution. A major focus of IF processing is to develop techniques that optimize resultant estimates, based on multisource inputs, usually in some statistical sense, to reduce the variance, uncertainty, or ambiguity.

This broad concept of IF is an important topic for a unified theoretical approach and therefore deserves its own label.[2]

IF Levels

Of the many possible ways to differentiate among types of IF functions, the JDL data fusion subpanel has become the most popular. This JDL model differentiates functions into fusion levels (depicted in figure 1) that provide a useful distinction among IF processes related to the assessment and refinement

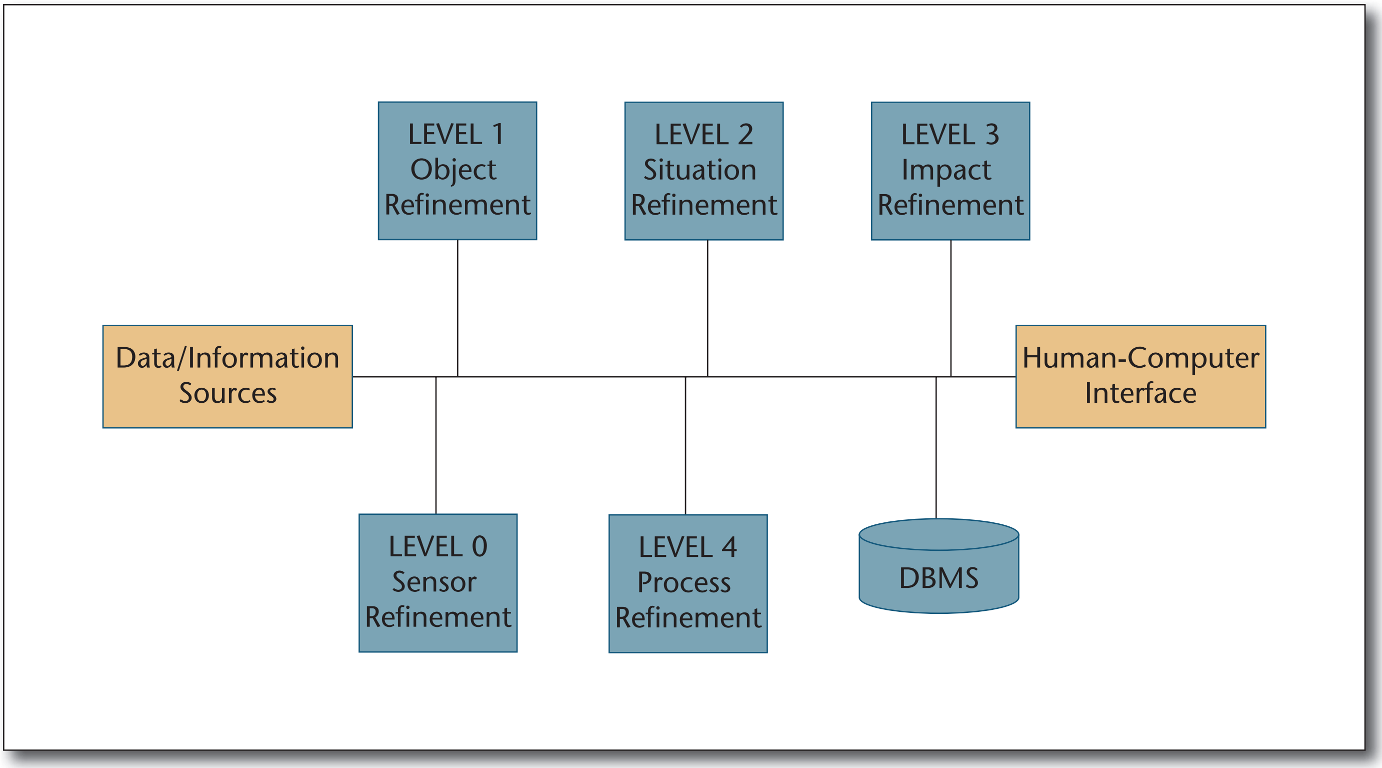

Figure 1. Classic JDL Data Fusion Process Model.

From Kessler et al. (1992).

of estimates for parameters of interest related to objects, situations, threats, and processes (Hall and Llinas 1997). Figure 1 depicts either a single IF node or the aggregate processing of a suite of IF nodes, each with similar structure. Figure 1 is strictly a discussion aid and not an architecture or processing diagram.

The JDL model (Hall and Llinas 1997) and subsequent revisions (Steinberg, Bowman, and White 1999; Llinas et al. 2004; Blasch et al. 2006), were proposed to provide a useful categorization representing logically different types of problems, generally (although not necessarily) solved by different techniques, and to maintain a degree of consistency with the mainstream of technical usage. Much of the discussion on IF models is drawn, either directly or in modified form, from the revisions to the JDL model by Steinberg, Bowman, and White (1999). We will use the following definitions:

> *Level 0 — Subobject data assessment.* Estimation and prediction of signal/object observable states on the basis of pixel/signal level data association and characterization. This is a new level added to the JDL process model, ascribed to the data sources. One clear example of this type of processing is the detection of blobs of interest in imagery, which signify an unknown entity of interest not having a semantic label.

Level 1 — Object assessment. Estimation and prediction of (ideally) single entity states on the basis of inferences from observations. A frequent example is the estimation of the kinematic properties and the identification or classification of single objects of interest in a situation.

Level 2 — Situation assessment. Estimation and prediction of entity states on the basis of inferred relations among entities. For a variety of applications, this level strives to develop situation assessments based on the collective dynamics among entities as well as the estimated relationships among them.

Level 3 — Impact assessment. Estimation and prediction of effects on situations of planned or estimated/predicted actions by participants (for example, assessing susceptibilities and vulnerabilities to estimated/predicted threat actions given one's own planned actions). In military applications, a major concern is threat estimation that involves estimating most likely or most damaging potential of an adversary.

Level 4 — Process refinement (an element of resource management). Adaptive data acquisition and processing to support mission objectives. A frequent focus of level 4 design and development is for a sensor management functionality, where observational strategies that improve the quality or reliability of fused estimates are developed.

Level 5 — User refinement (an element of knowledge management). Adaptive determination and retrieval of data to support cognitive decision making and action; this is a new level added to the JDL process model that extends the functions of the human-computer interface. For many applications, the human can provide a number of useful functions to the operations of a fusion process, especially to include aspects of priority or value or of human intent.

IF and Context

One overarching assertion made in this article is that IF is fundamentally an estimation process, specifically an autonomous or largely autonomous estimation process. The central focus in this article is on algorithmic or otherwise autonomous processes, although the insertion of human intelligence with the automated systems is not disallowed. These automated processes are enabled in software and thus are vulnerable to the bad-data-in/bad-data-out constraint. IF functional capabilities, especially at the highest levels, are developed to estimate some aspect of a real world of interest to some user, from which those capabilities are bounded by a number of factors, such as the quality of the available data or information used to form the estimates,[3] as well as the complexity of the world being observed. Gauging the complexity of any world situation includes the degree to which that world is hospitable to, or cooperative with, the observational or other information-providing mechanisms on which the IF estimation processes depend. In adversarial domains such as defense and security domains, including adversarial business postures in the business world, there can be purposeful actions to deceive or corrupt any IF process; such environments lead to the endless exchange of actions intended to overcome each step of an adversary, in an endless game (countermeasures on countermeasures). IF researchers and developers hope to enable an IF capability that aids human sensemaking (Klein, Moon, and Hoffman 2006; Pirolli and Card 2005) or, in general, reasoning. Thus, it is not unexpected that IF process designs strive to mimic, support, or augment human reasoning.

Reasoning is generally seen as a means to improve knowledge and make better decisions, although empirical evidence quantifies a number of error types, biases, and inefficiencies in human reasoning. The literature distinguishes between reasoning and inferencing; for example, see Steinberg and Bowman (2009). Mercier and Sperber (2011) define inferencing as follows:

> Inference (as the term is most commonly understood in psychology) is the production of new mental representations on the basis of previously held representations. … Reasoning, as commonly understood, refers to a very special form of inference at the conceptual level, where not only is a new mental representation (or conclusion) consciously produced, but the previously held representations (or premises) that warrant it are also consciously entertained.

Continuing, Mercier and Sperber (2011) say that what characterizes reasoning is that it includes and is distinguished by the notion of a formed argument, its purpose oriented to persuasion. In the argument approach, the premises that are reasoned over are seen as providing reasons to accept or reject a conclusion. Thus, what characterizes reasoning from inferencing is the awareness not just of a conclusion but of an argument that justifies accepting that conclusion. Furthermore, the distinction between reasoning and inference is based on a premise, most definitions labeling a premise as a statement that is assumed to be true and from which a conclusion can be drawn.[4] So, premises are statements or assertions that support the forming of a conclusion. Utilizing a statement to support a conclusion in high-level IF (HLIF) processes requires the use of contextual information together with observed evidence.

In broad terms, contextual information is information that surrounds a situation of interest in the world — information that aids in understanding the (estimated) situation and also in reacting to the situation, if reaction is required. Most problems addressed by IF systems are related to interpretation and development of meaning from multisource data. Therefore, there is usually some focal data collected to develop such understanding. For instance, in a surveillance application, these are sensor data and possibly human-based observational data. Through analysis, these data can support the formation of focal premises (statements about some aspect of the condition or situation of interest). To the extent that separate contextual data or information is available, these too can be analyzed to form additional premises — propositions that we call *contextual premises* – that, together with the focal premises, can lead to the formation of an argument, that is, a conclusion traceable to the foundations of the joint set of these premises.

Context exploitation can provide benefits in IF by establishing expectations of world states, explaining and constraining received data, and resolving ambiguous interpretations (Rogova and Steinberg 2016). Context allows for improving the associability between problem-space knowledge, derived models, and observational data. These a priori (that is, mathematical) models can better explain data by exploiting the semantics provided by contextual information (Gómez-Romero et al. 2015; Snidaro et al. 2013).

The typically rich semantics associated with contextual data can also help build systems that can explain their reasoning process. Hence, contextual

	First Wave	Second Wave	Third Wave
Years	1960s–1980s	1980s–2010s	2010s–
Technology	Expert Systems	Machine Learning	Deep Learning
Algorithms	Logical rules	Statistical methods	Statistical methods
Expert knowledge	Expert knowledge → Rules	Expert knowledge → Model, Features	Expert knowledge → Model
Learning		Data → Parameters	Data → Parameters
Algorithm application	Rules → Data	Model → Data	Model → Data
Handling of uncertainty	NO	YES	YES
Abstraction	NO	NO	YES
Interpretable	YES	NO	NO

Table 1. Major Phases in AI History.

information can enhance AI systems for explanations of decisions.

Context and AI

The goal of AI research is to bestow on machines the ability to solve problems by mimicking human intelligence. The inherent power of modern electronic devices to process enormous quantities of data and information would therefore be augmented with humanlike intelligence.

AI research spans many areas, among which perception and cognition play a pivotal role in developing systems that need to be aware of the surrounding environment and have the capability to assess and adapt to it. The term cognition covers fundamentals concepts such as reasoning and learning.

As discussed in a recent overview on the state of AI research (Deng 2018), since the 1960s, the development of algorithms simulating human intelligence has gone through three major phases or waves. Table 1 summarizes the technology and capabilities characterizing the solutions that have been developed in these three phases, mostly regarding attempts made over the years to address aspects relating to human cognition.

The first wave defines expert systems that tried to simulate human reasoning to solve specific problems in very narrow domains (such as medical diagnoses). The intelligence of the software is mainly encoded in a set of logic rules that make up a rule base able to process (mostly in a deductive manner) the queries of the user according to the known facts contained in a knowledge base.

The development of the rule base was a painstaking and laborious process that had to be carried out to encode, in the form of logic rules, using the knowledge of a human expert (such as a medic). Expert systems were unable to learn from data, were brittle at handling uncertainty, and had no ability to abstract new concepts (that is, develop more general rules) from data. The major advantage of these systems was their interpretability, that is, the possibility to trace exactly, in a humanly understandable way, the chain of reasoning that led to certain conclusions. The trace usually meant looking at the sequence of logic rules firing on a given input. Given their overall simplicity and amenable property of relying on an interpretable reasoning mechanism, expert systems are still used today, especially considering the development that has been added to them over the years to mitigate their shortcomings (for example, adding uncertainty management via fuzzy logic to subsume both predicate logic and probability theory [Zadeh 1983]).

The second wave marks the advent of machine learning, such as statistical learning methods that are able to solve a given problem (for example, a classification problem) by learning from a number of examples and their corresponding expected output. In statistical learning, there are no explicit rules being provided to the system, and the effort of the developer is in choosing the most convenient model and parameters, describing the problem in terms of a number of characteristics or features, and having the system learn from examples according to a learning algorithm (for example, backpropagation). The learning process discerns the internal parameters of the statistical model so that the system is flexible to variations of the input features and provides a confidence level on its output. Therefore, on being presented a new input, the system is able to provide a classification result accompanied by a degree of confidence. These statistical learning systems can handle never-seen-before and noisy inputs. However, these methods, especially neural networks, operate like black boxes, making it is impossible to attribute semantics to the values of the parameters learned during the training phase. Statistical learning algorithms such as neural networks are, in general, not interpretable, making it very difficult to explain or visualize the reasoning behind the results.

In the past few years, the interest in AI has grown considerably thanks to the advent of what Deng (2018) describes as the third wave, that is, a new paradigm called deep-structured machine learning or deep learning. The approach involves the use of deep neural networks comprising several hidden layers. These models, trained on very large data sets with sufficient computing power (boosted by the availability of specific architectures, such as graphics processing units), have far surpassed the performance of the shallow neural network models typical of the second wave. For example, speech and image recognition are tasks for which the deep paradigm has been particularly successful and is starting to show human-level performance.

The success of deep neural networks mostly relies on their ability to create abstractions from the observed training data. The abstraction ability gives them great generalization capabilities that allow them to perform better in the presence of noisy input data. The abstraction capability marks another significant difference with respect to the previous generation. For example, the careful features engineering process is no longer required in tasks like image classification, where the raw data (for example, image pixels) are directly fed to the network. Hence, deep networks can take as input raw signals, or a subsampled version of them, and directly learn the most significant features through different layers of abstraction for the classification task at hand. These models require the developer to focus only on the selection of the most appropriate model, its configuration, and choice of hyperparameters (for example, learning rate, number of epochs). One shortcoming of the deep neural network paradigm, however, inherited from previous-generation shallow models, is again the nonexplainability of the results. Once again, the networks operate as a black box on the inputs and transform them into an output based on the weights and parameters of the network learned during the training phase. However, active research is being done on the subject with some very recent promising results, as in the case of Shapley additive explanations (Lundberg and Lee 2017).

In all phases of AI, CI has played a significant role, even when not explicitly recognized. Expert systems were developed for specific problems in specific contexts, and the failure to correctly represent the contextual factors was often one of the causes of their malfunctioning. The systems were not portable or easily transferrable to conditions different from those for which they had been developed. In a sense, all the contextual knowledge that the expert had provided was so hard coded and built-in that any attempt to test the system even slightly outside of that context produced unreliable outputs. In the second phase, context was exploited only in recent years (Snidaro, Garcia, and Llinas 2015) as a way to boost the performance of certain tasks. For example, object classification in images was shown to benefit from considering the pixels surrounding the object of interest. The surrounding pixels provided additional information that allowed computations to better recognize a certain object given its background and context. More direct approaches have been taken with the explicit design of additional nodes or layers to store contextual information on deep learning, as by Delcroix et al. (2016).

The next wave of AI will need to put more focus on how to reason over data by considering contextual factors and by incorporating contextual models over time in the learning process. The ability to understand and exploit contextual elements will drive the development of future intelligent systems toward explainable results.

	Artificial Intelligence	Information Fusion	JDL Level
Perception	Machine perception	Low-level information fusion	0, 1
Cognition	Reasoning	High-level information fusion	2, 3, 5
Adaptation	Learning	Process refinement	4

Table 2. Dimensions and Key Technologies for Advanced AI and IF Systems.

Whereas the second wave of research emphasized inductive learning of statistical distributions or discrimination functions, a more general approach is required for the future of AI systems to provide explainable decisions, which should be based on conceptual representation and allow automatic adaptation to the conditions of the domain. For instance, a model focused on concepts and concept learning is the model of rational rules (Goodman et al. 2008; Lake, Salakhutdinov, and Tenebaum 2015), which combines the inferential power of Bayesian induction with the representational power of mathematical logic and generative grammars for concept generalization. Similarly, Markov logic networks (see the work of Snidaro, Visentini, and Bryan [2015] for an application for maritime SAW) integrate logic representation with models of uncertainty. These are probabilistic models developed on Markov networks with first-order logic to enable inferences under uncertainty. In these cases, the possibility of using a symbolic representation of the concepts learned allows the system both to generalize and to adapt to specific conditions for each domain.

A fundamental challenge for next-generation AI systems, moving toward even more intelligent capabilities, will be the ability to adapt to contextual conditions, as discussed in the following sections.

Fusion and AI

The next wave of AI challenges should emphasize explainable models instead of black-box paradigms. The key is providing explanations to users, especially in the mission-critical applications required to facilitate human-machine interaction. In this regard, there is a parallelism between the goals of the AI and IF research communities in the current strategies to achieve progress in several dimensions, identified in table 2. There are distinguishable layers of functionalities needed to develop intelligent systems: perception is the connection with the physical world, access to data about entities to be represented and characterized; cognition is the process of understanding the perceived world, where situation is the key concept for this layer, as we will explain below; and adaptation is the strategy to optimize the performance of perception and/or cognition to better reach a certain goal.

In the development of IF systems for decision support (where high-level fusion processes are involved), the representation of the world according to a certain model should explain decisions made. The HLIF model (for example, see the work of Blasch, Bosse, and Lambert [2012]) can be directly aligned with Endsley's SAW model (Endsley 1995), as used by the IF community to characterize the mental process carried out by experts to understand situations. The SAW model could provide inspiration toward explainable results of future AI systems. Per Endsley,

> Situation awareness is the perception of elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future.

The SAW model is basically organized into three levels that correspond with the JDL data fusion model described earlier, where cognition and comprehension are similar, while projection is a subset of adaptation. Note that the IF community has furthered the concept of a human computer interface of the JDL model by developing the methods for level 5 user refinement utilizing the SAW model concepts, which allow the user to input cognition (Blasch, Bosse, and Lambert 2012).

The alignment of the JDL and the SAW models supports perception of situational elements, such as objects and events, comprehension as understanding the meaning of situational elements, and projection as awareness of the future evolution of the situation. Perception involves the detection of entities and events by sensors and data fusion techniques. This level can be correlated with machine perception, the first stage of intelligent systems to take input data from real-world scenes. Comprehension is focused on situation understanding by analysis of the relationships among the detected entities and events. This process usually requires a formal representation to detect and instantiate the relationships among the detected entities and requires reasoning techniques able to deal with their evolution in time. A process of projection estimates the evolution of a situation by using inference processes This level benefits from machine learning by intelligent systems to forecast situations based on models.

Therefore, the two main aspects to developing HLIF systems are representation of situations and reasoning about their expected outcomes. Regarding representation strategies, Strang and Linnhoff-Popien (2004) considered several major types of models, of which three are most applicable to IF: key-value models, ontology-based models, and logic-based models (for example, the ontology web language [OWL]). A formal approach for reasoning about situations (Powell et al. 2006) is the use of ontologies to represent objects and relationships as well as their evolution over time. An ontology is a specification of concepts and relationships among entities that can exist in a given setting (Kokar, Matheus, and Baclawski 2009). In the ideal case, this description can express in a reasonable way the evolution of objects and their relationships (Gómez-Romero, Bobillo, and Delgado 2011). Recent researchers in IF have recognized the advantages of cognitive situation models, pointing out the importance of formal context knowledge to achieve scene understanding. Steinberg and Bowman (2009) emphasized the importance of context knowledge, for instance, when visual inputs are to be interpreted (Gómez-Romero et al. 2012). Sycara et al. (2009) state that parts of the context are the significant features or the history of a situation that influence the features of other situations, as well as the expectations on what is to be observed and the interpretation of what has been observed.

With respect to reasoning, ontological representations provide a logically robust classification of the entities in its domain, which can be used to analyze and derive relationships among instances of entities in certain situations. Ontologies formally describe the core concepts, the facts, and their interrelationships, which can be used to build realistic representations of context (Gómez-Romero et al. 2010), and they can be extended with ad hoc logic-based models for knowledge-intensive applications. Ontologies represent context as facts and information inferred from rules to allow for the development of more sophisticated representations and reasoning procedures. Ontologies must use highly expressive, logic-based knowledge models to describe a domain with an automatically processable language. Description logics (DLs) are a family of logics used to represent this structured knowledge (Baader, Horrocks, and Sattler 2005) and have proved to be suitable ontology languages. Ontology languages are usually equivalent to a decidable DL, as occurs in the case of OWL2. DLs are categorized by level, according to the allowed expressions in them. In general, the more expressive a language, the greater the computational complexity required for inferences. The situation theory ontology, modeled in OWL, is the result of extensive research in situation theory and provides a formal language for representing relationships between situational elements.

With respect to inference and reasoning processes for situation and impact refinement (JDL level 2 or 3), some initiatives such as the rule interchange format search formats for the interchange of rules in rule-based systems to create a common interchange format for different rule languages and inference engines. The basic reasoning task regarding ontology concepts is satisfiability, which implies there are no contradictions with the rest of the stored knowledge. Another important task is the concept of subsumption, which occurs when a concept is more general than another concept. So, a typical inference done over ontology instances is to test whether there are not contradictions in the axioms in the knowledge base. The DL inference engines execute these types of tasks over ontologies built with OWL and/or OWL 2 languages, such as Pellet (Sirin et al. 2007) and RACER (Haarslev and Moller, 2001) as two representative examples of DL reasoners. The computational complexity of the reasoning procedures depends on the expressivity of the language considered. Rule-based reasoning is not directly supported by OWL, but several extensions have been proposed. One of the most interesting is the semantic web rule language (Sirin et al. 2007), an extension of OWL, which includes deductive inference within OWL ontologies.

Therefore, context adaptation is one of the top challenges identified for the next wave of AI research, a goal aligned with the IF research community to develop advanced systems able to represent situations and reason about them. Developing adaptive processes (that is, JDL level 4 processes according to the model presented at the beginning of this article) requires an explicit representation of situations and contextual premises in the pathway to reason about the process and to trigger appropriate adaptation mechanisms. Likewise, user refinement (that is, level 5 fusion) requires systems with human explainable reasoning for interpretability.

Perspectives on Context for IF and AI

A cornerstone for new advances in the AI field is the development of interpretable models, a fundamental aspect to explaining observations according to the knowledge acquired and explaining the decisions made depending on situations. These interpretable models should take into account the available contextual data (static and dynamic), which implies adaptation to the conditions of the problems addressed. In the same way, context is fundamental to deriving adaptive IF systems. This section analyzes the role of context in the research of both communities.

Context Representation and Semantics: Joint Perspective on Context from AI and IF

Historically, AI has always dealt with the notion of context, which has now become the focus of challenges in the next wave of intelligent systems. A pioneering approximation to context formalization in the field of AI is by McCarthy (1993), who proposed extensions of logic relations to explicitly include context. Giunchiglia (1992) defines an analogous framework in which the context is a subset of the complete state of an entity used to solve a task. These approaches have been investigated to address context modeling with ontologies in the semantic web. From that context modeling perspective, and directly related to reasoning, context representation in AI can be related to maintaining SAW for understanding an evolving situation (Kokar, Matheus, and Baclawski 2009). Another aspect, and a key challenge for AI researchers, is the development of commonsense reasoning, identified as a fundamental problem in complex real-world tasks such as computer vision, robotic manipulation, and temporal planning (Davis and Marcus 2015). A clear example is autonomous systems working in real environments with unpredictable situations, where the systems must react appropriately to unanticipated events and so need a type of reasoning that allows them to effectively react and avoid nonsense inferences. A possible strategy to implement common sense in intelligent systems is to use taxonomic reasoning with a corpus of events, actions, and expected changes. A taxonomy is a collection of categories and individuals and the relationships between them, and it can be assimilated within the ontologies representing the situations indicated earlier. Thus, developing models integrating the known taxonomies of a certain domain is a way to build explainable models that are sensitive to context.

Building explainable models is also directly related to the paradigm of general intelligence and cognitive architectures. Toward upgrading machine learning to human-level learning and problem-solving abilities, a goal known as artificial general intelligence has been chased by several AI researchers (Voss 2007). The basic assumption of this paradigm is that the study of human cognition can generate computational models of human behavior.

Cognitive architectures are based on the model of interactive learning, adapting to changing circumstances, abstraction, and reuse of knowledge and skills, as well as reasoning and language understanding. The path to artificial general intelligence with cognitive architectures is based on the active search and accumulation of knowledge from various (or changing) environments, with the same basic mechanism of SAW: perceptive observation of features from the world, cognitive representation of the knowledge obtained, and adaptive output/actuation mechanisms.

In this way, cognitive architectures (Langley 2006; Langley 2012) focus on autonomous concept building. The notion of a cognitive architecture is associated with the approach of agent design in terms of organizing short-term and long-term memories that store the content and control of its behavior. One or more long-term memories contain knowledge and procedures (skills), along with short-term memories that store the agent's beliefs and goals. In addition, functional processes operate on these structures, such as learning mechanisms. The cognitive paradigm corresponds well with the framework explored by community research of autonomous agents, known as belief-desire-intention (Rao and Georgeff 1995), where a decision model for logic deliberation is based on available perceptions and knowledge represented as a tree of possible worlds, each one named as a situation. A balance between reactive and planned behavior is achieved by committing agents to plans but periodically reconsidering the plans according to the perceptions of dynamic situations. A high-level conceptual representation of the world includes knowledge about a domain, but, at the same time, the plan building mechanism allows reactivity to the context of situations.

From the perspective of IF design, as mentioned, computational SAW has been supported by extensive research in the theory of SAW and technologies of the Semantic Web (ontologies and inference engines). The resource description framework and OWL, both W3C recommendations, are expressive graph-based knowledge representation languages that provide the ability to effectively represent relationships. A strategy known as semantic enhancement of data is based on ontologies and links built to provide better retrieval and integration. Further, ontologies for semantic enhancement enable analytics tools to see through the inconsistencies and redundancies in data. For example, in the decision support domain, new search techniques are needed for human analysts to deal with heterogeneous data scattered in different corpora with multiple terms and data models (Smith et al. 2011).

Perspectives on Context for Fusion Adaptation

With the goal of developing an automated IF system exploiting contextual information, it must be considered that context can be dynamic and changing at perhaps a different rate than how the problem's variables evolve. A context middleware concept has been proposed (Snidaro et al. 2015) to retrieve and provide timely CI to the fusion processes affected by this information. Additionally, it includes assessing relevance for deciding whether the details of the data and information surrounding an area of interest are pertinent and, following the terminology introduced in the section on context and AI, what contextual premises will influence the arguments about the situation that the fusion system is trying to develop. Thus, a fundamental decision in the life cycle of the IF system of interest, when this question about relevance will be asked, leads to the two main alternatives for context integration as described by Rogova and Snidaro (2018).

The a priori CI exploitation framework takes advantage of knowing the arguments or estimates the IF system is designed to perform. In the a priori case, at design time, one knows the goal arguments that the system will be trying to develop, taking into account the physical, social, and perhaps other contextual information that could be available, from which relevant premises could be formed.[4] With goal arguments, the a priori label for the framework formed at design time for the exploitation of contextual information attempts to account for the effects of contextual premises on goal argument formation (such as situational estimation). Integrating the focal, sensor-data-derived premises with the context-information-derived premises into a combined argument framework is a separate design problem, with the associated ease or difficulty involved in integrating contextual premises into a fusion system design or into any AI algorithm design. In the a priori framework, the system assumes designers know beforehand what CI-type information is relevant to the intended goal arguments of the fusion process.

In a different scenario, an a posteriori CI exploitation framework can be proposed for the case when all relevant CI may not be known at system or algorithm design time and may have to be searched and discovered at run time as a function of the current goal argument or situation estimate and evolving mission objectives. Some CI may not be of a type that allows easy integration into the system or algorithm designs at design time and so may not be easily integrated into the goal argument or situation estimation process. In this case, at least part of the a posteriori exploitation task should be to check the consistency of a current fused argument with the newly discovered (and relevant) CI, adding explanatory aspects to the declared hypothesis. That is, if the current argument or hypothesis is also consistent with the new, additional CI, that argument should be tagged as such, indicating that it is a stronger hypothesis.

Context Adaptation Architecture

A middleware-based context-enhanced IF architecture design was proposed by (Snidaro et al. 2015). The basic mechanism proposed follows a query-response middleware interface between fusion processes and CI — where the selected relevant CI from available sources is available according to the values inferred and hypotheses proposed by fusion processes. Two basic elements can be identified in both sides.

At the context side, the middleware is responsible for collecting and updating context knowledge and making it usable by fusion processes. At the fusion side, the fusion adaptation logic uses the contextual inputs, so all processes and modules need to be described in terms of context input and interconnections to apply the adaptation.

At the fusion process side, the CI reported as relevant in the response to a query is exploited to improve the performance following an adaptive strategy. The adaptation mechanisms are based on alternative models when they can be selected according to context (such as on/off road motion models); measurable impact of applicable models, sets of parameters, and algorithms; and applicable rules to drive the fusion processes, such as constraints, applicable models, and hypotheses.

Context can be exploited in any of these functions, independently or jointly, to adapt the process to the available knowledge. The adaptation of the fusion system is motivated and guided by two sources of change: the purpose of the system (desired data products and their features) and the quality/performance of the fusion processes. The purpose is the most ambitious, an online reconfiguration logic of the problem space originated either by human operators or by intelligent external entities controlling the fusion system. The quality and performance of the fusion process can be evaluated locally at each fusion node through the analysis of intermediate indicators or results. The quality of the fusion products, thus, is the principal driver for the adaptation process.

The adaptation architecture facilitates adaptation processes through the scheme indicated in figure 2 at two levels: adapting the configuration of individual components as sensors or fusion nodes and adapting the structure of the solution. As shown, context-based adaptation takes place in several places, as follows.

A sensor management module has a global view of the sensor set, geographical disposition, and the data needed by the fusion process. It is in charge of managing parameters such as deployment or space-time scheduling of sensors that best satisfy data needs.

Individual sensor modules have adaptive logics responsible for adjusting the internal parameters of their sensors. These parameters respond to two needs: improve overall quality of the data and maximize certain features of the obtained data (for example, a radar can be configured to improve refresh rate, or to discern the shape/size of a target).

The overall quality control process module is in charge of evaluating the combined fusion product to identify performance impacts and opportunities at an interlevel view. This information is used to control two aspects of the fusion process: orient each level of the fusion system (this can affect its relationship with other levels or modify its goals) and request changes in the sensor set (scheduling and management).

Internodal adaptive logic modules evaluate the combined fusion products generated by fusion nodes at a certain level. They are used to determine how to control two aspects of the fusion process: individual parameterization/algorithm selection at each fusion node and relationship between fusion nodes.

At the top of level 4 processes, the fusion adaptation block is the access point to the available

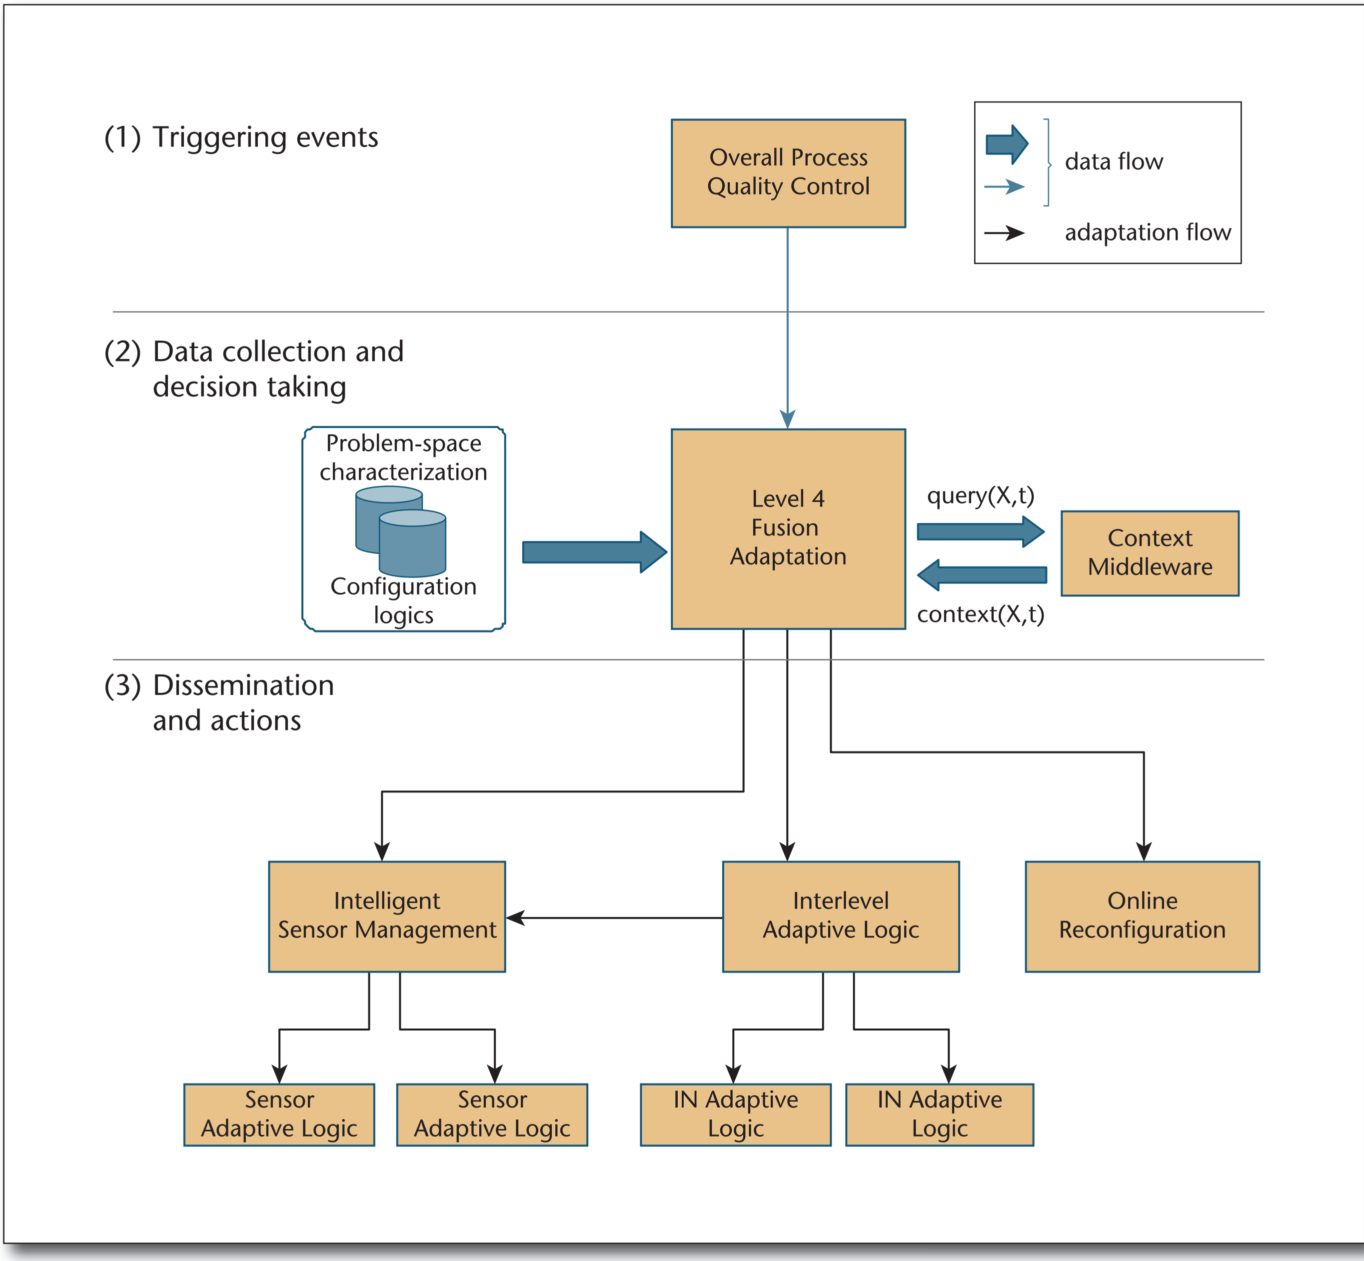

Figure 2. Data Flow View of the Context-Based Adaptation Process.

contextual information and can also access a description of the problem space. Context is used by adaptation management to identify needs or desired changes in the flow control and sensor scheduling. These needs and changes are disseminated together with relevant context data to the sensor management module and the interlevel adaptive logic.

Once the system has the relevant context data, the interlevel logic can define the actions to be taken to improve the process, both inter- and intralevel. Context-based sensor management receives requests for data, data features, and sensor time-space pointing from other components of the IF system. It is in charge of managing the sensor sets to satisfy requests as well as possible.

A sensor adaptive logic knows the relationship between the configuration parameters and the effect on the data captured by the sensor. The logic knows how to change sensor parameterization (capabilities/features of the sensor: frequency, resolution, type of information) to obtain the desired effect. The reasoning process of the sensor adaptive logic is expected to be fairly simple. Referring to the previous example, a radar sensor adaptation process could ask for better resolution of target sizes or shapes in a certain region, which is translated into a different pulse width or antenna rotational speed.

Finally, the internal configuration logics are used to provide the adaptation module with the information it needs to plan the adaptive actions. The mission of the online reconfiguration module is to change the internal structure of the fusion solution to satisfy the goals or better respond to external circumstances (context). Therefore, the relevant context and changes in the goals are the two sources used to decide actions over the structure of the fusion solution — activation/deactivation of fusion nodes or changes in the control flow.

Conclusions

The article identifies handling of CI and an appropriate situation representation as key elements to develop adaptive and explainable AI systems. Although AI is certainly present in current applications and has become a popular area of computer science, in part because of business investments and media coverage of AI success in games, some open challenges have been discussed to motivate progress in this area, called the next milestone or the third wave of AI. Whereas previous milestones have been characterized by manually crafted knowledge (expert systems) and statistical machine learning, the open challenge is now to learn paradigms with explainable models and context adaptation to gain generalization capability. Explainable AI can benefit from IF developments, context enhancement, and decision support systems.

AI must address many problems that are current in the area of IF: perception and object recognition from sensory data or SAW. A basic challenge identified for both AI and IF future systems is understanding context, the ability to represent and relate how relevant the context is to the inference problems addressed, and mechanisms to adapt the inference processes to this context. In this parallelism, the challenges include perception, reasoning, and adaptation toward deploying AI and IF systems to support knowledge representation and situation understanding.

The preeminent challenge is context adaptation common to both AI and IF research. A basic objective is the capability to learn interpretable models from contextual data to bind observations with knowledge and use the semantics provided by context. Current strategies for situation assessment and HLIF have been reviewed and related with the ongoing research in AI branches of computer science (compositional learning models, cognitive architectures, common-sense reasoning, and rational agents). A general approach for AI systems should be based on conceptual representation to allow automatic adaptation to domain conditions. Future breakthroughs will stem from an architecture to represent, access (middleware), and exploit the context in IF processes that provide context-enhanced systems.

Notes

1. See www.darpa.mil/program/explainable-artificial-intelligence.

2. Terms like information integration have been preferred by some to connote greater generality than earlier, narrower definitions of data fusion (and, perhaps, for distance from old data fusion approaches and programs), but such manipulations do not contribute toward better representation or understanding.

3. Quality is another term for which it is difficult to achieve consensus; data quality has been written about for general applications (Sycara et al. 2009) as well as for IF-specific applications (Zadeh 1983). We address the quality issue as part of the general content of this article in various ways.

4. A definition of a priori is “formed or conceived beforehand,” and here we mean it in the sense of incorporating contextual aspects in a design at the outset.

References

Baader, F.; Horrocks, I.; and Sattler, U. 2005. Description Logics as Ontology Languages for the Semantic Web. In *Mechanizing Mathematical Reasoning*, edited by D. Hutter and W. Stephan, 228–48. Vol. 2605. Lecture Notes in Computer Science. Berlin: Springer. doi.org/10.1007/978-3-540-32254-2_14

Biran, O., and Cotton, C. V. 2017. Explanation and Justification in Machine Learning: A Survey. Paper presented at the International Joint Conference on Artificial Intelligence 2017 Workshop on Explainable Artificial Intelligence. Melbourne, Australia, August 20.

Blasch, E.; Kadar, I.; Salerno, J.; Kokar, M. M.; Das, S.; Powell, G. M.; Corkill, D. D.; and Ruspini, E. H. 2006. Issues and Challenges in Situation Assessment (Level 2 Fusion). *Journal of Advances in Information Fusion* 1(2): 122–39.

Blasch, E. P.; Bosse, E.; and Lambert, D. A. 2012. *High-Level Information Fusion Management and Systems Design*. Norwood, MA: Artech House.

Davis, E., and Marcus, G. 2015. Commonsense Reasoning and Commonsense Knowledge in Artificial Intelligence. *Communications of the ACM* 58(9): 92–103. doi.org/10.1145/2701413.

Delcroix, M.; Kinoshita, K; Yu, C.; Ogawa, A.; Yoshioka, T.; and Nakatani, T. 2016. Context Adaptive Deep Neural Networks for Fast Acoustic Model Adaptation in Noisy Conditions. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*. Picataway, NJ: IEEE.

Deng, L. 2018. Artificial Intelligence in the Rising Wave of Deep Learning: The Historical Path and Future Outlook. *IEEE Signal Processing Magazine* 35(1): 180, 173–7. doi.org/10.1109/MSP.2017.2762725

Endsley, M. R. 1995. Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors Journal* 37(1): 32–64. doi.org/10.1518/001872095779049543.

Giunchiglia, F. 1992. Contextual Reasoning. Epistemologia 345: 345–364.

Gómez-Romero, J.; Bobillo, F.; and Delgado, M. 2011 Context Representation and Reasoning with Formal Ontologies Paper presented at the Activity Context Representation Workshop in the 25th AAAI Conference, San Francisco, CA, August 7–8.

Gómez-Romero, J.; García, J.; Kandefer, M.; Llinas, J.; Molina, J. M.; Patricio, M. A.; Prentice, M.; and Shapiro, S. C. 2010. Strategies and Techniques for Use and Exploitation of Contextual Information in High-Level Fusion Architectures. In *Proceedings of the 13th International Conference on Information Fusion*. Piscataway, NJ: IEEE. doi.org/10.1109/ICIF.2010.5711859

Gómez-Romero, J.; García, J.; Patricio, M. A.; Molina, J. M.; and Llinas, J. 2011. High-Level Information Fusion in Visual Sensor Networks. In *Information Processing in Wireless Sensor Networks: Technology, Trends and Applications*, edited by K.-L. Ang and K.-P. Seng, 177–223. Hershey, PA: IGI Global.

Gómez-Romero, J.; Serrano, M. A.; García, J.; Molina, J. M.; and Rogova, G. 2015. Context-Based Multi-Level Information Fusion for Harbor Surveillance. *Information Fusion* 21: 173–86. doi.org/10.1016/j.inffus.2014.01.011.

Gómez-Romero, J.; Serrano, M. A.; Patricio, M. A.; García, J.; and Molina, J. M. 2012. Context-Based Scene Recognition from Visual Data in Smart Homes: An Information Fusion Approach. *Personal and Ubiquitous Computing* 16(7): 835–57. doi.org/10.1007/s00779-011-0450-9.

Goodman, N. D.; Tenenbaum, J. B.; Feldman, J.; and Griffiths, T. L. 2008. A Rational Analysis of Rule-Based Concept Learning. *Cognitive Science* 32(1): 108–54. doi.org/10.1080/03640210701802071.

Haarslev, V., and Moller, R. 2001. Description of the RACER System and Its Applications. In *Proceedings of the International Workshop on Description Logics (DL-2001),* Stanford, CA, August 1–3, pp. 132–141.

Hall, D. L., and Llinas, J. 1997. An Introduction to Multisensor Data Fusion. *Proceedings of the IEEE* 85(1): 6–23. doi.org/10.1109/5.554205.

Kessler, O.; Askin, K.; Beck, N.; Buede, D.; Hall, D.; Llinas, J.; Lynch, J.; and White, F. 1992. Functional Description of the Data Fusion Process. Technical Report. Warminster, PA: Naval Air Development Center.

Klein, G.; Moon, B.; and Hoffman, R. R. 2006. Making Sense of Sensemaking 2: A Macrocognitive Model. *IEEE Intelligent Systems* 21(5): 88–92. doi.org/10.1109/MIS.2006.100.

Kokar, M.; Matheus, C.; and Baclawski, K. 2009. Ontology-Based Situation Awareness. *Information Fusion* 10(1): 83–98. doi.org/10.1016/j.inffus.2007.01.004.

Lake, B. M.; Salakhutdinov, R.; and Tenenbaum, J. B. 2015. Human-Level Concept Learning Through Probabilistic Program Induction. *Science* 350(6266): 1332–8. doi.org/10.1126/science.aab3050

Langley, P. 2006. Cognitive Architectures and General Intelligent Systems. *AI Magazine* 27(2): 33–44. doi.org/10.1609/aimag.v27i2.1878

Langley, P. 2012. The Cognitive Systems Paradigm. *Advances in Cognitive Systems* 1: 3–13.

Little, E. G., and Rogova, G. L. 2009. Designing Ontologies for Higher Level Fusion. *Information Fusion* 10(1): 70–82. doi.org/10.1016/j.inffus.2008.05.006

Llinas, J.; Bowman, C. L.; Rogova, G. L.; Steinberg, N.; Waltz, E. L.; and White, F. E. 2004. Revisiting the JDL Data Fusion Model II. In *Proceedings of the Seventh International Conference on Information Fusion*. Stockholm, Sweden: FOI.

Lundberg, S. M., and Lee, S. I. 2017. A Unified Approach to Interpreting Model Predictions. *Advances in Neural Information Processing Systems* 30: 4765–74.

McCarthy, J. 1993. Notes on Formalizing Context. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence*. San Francisco: Morgan Kaufmann Publishers.

Mercier, H., and Sperber, D. 2011. Why Do Humans Reason? Arguments for an Argumentative Theory. *Behavioral and Brain Sciences* 34(2): 57–74.

Pirolli, P., and Card, S. 2005. The Sensemaking Process and Leverage Points for Analyst Technology as Identified through Cognitive Task Analysis. In *Proceedings of the International Conference on Intelligence Analysis*. McLean, VA: MITRE Corporation.

Powell, G.; Matheus, C.; Kokar, M.; and Lorenz, D. 2006. Understanding the role of context in the interpretation of complex battlespace intelligence In *Proceedings of the 9th International Conference on Information Fusion (Fusion 2006),* doi.org/10.1109/ICIF.2006.301719

Rao, S., and Georgeff, M. P. 1995. BDI Agents: From Theory to Practice. In *Proceedings of the 1st International Conference on Multiagent Systems*. Cambridge, MA: MIT Press.

Rogova, G., and Snidaro, L. 2018. Considerations of Context and Quality in Information Fusion. In *Proceedings of the 21st International Conference on Information Fusion*. Piscataway, NJ: IEEE. doi.org/10.23919/ICIF.2018.8455755

Rogova, G. L., and Steinberg, N. 2016. Formalization of Context for Information Fusion. In *Context-Enhanced Information Fusion, Advances in Computer Vision and Pattern Recognition*, edited by L. Snidaro, J. García, J. Llinas, and E. Blasch, 27–43. Berlin: Springer.

Sirin, E.; Parsia, B.; Cuenca Grau, B.; Kalyanpur, A.; and Katz, Y. 2007. Pellet: A Practical OWL-DL Reasoner. *Journal of Web Semantics* 5(2): 51–3. doi.org/10.1016/j.websem.2007.03.004

Smith, B.; Malyuta, T.; Salmen, D.; Hansen, A.; and Cronen, S. 2011. Ontology for the Intelligence Analyst. Integration of Intelligence Data through Semantic Enhancement, In *Proceedings of the Conference on Semantic Technology in Intelligence, Defense and Security*.CEUR-WS.

Snidaro, K.; Visentini, I.; Llinas, J.; and Foresti, G. L. 2013. Context in Fusion: Some Considerations in a JDL Perspective. In *Proceedings of the 16th International Conference on Information Fusion*. Piscataway, NJ: IEEE.

Snidaro, L.; García, J.; and Llinas, J. 2015. Context-Based Information Fusion: A Survey and Discussion. *Information Fusion* 25: 16–31. doi.org/10.1016/j.inffus.2015.01.002

Snidaro, L., J. García, J. Llinas, and E. Blasch, editors. 2016. *Context-Enhanced Information Fusion: Boosting Real-World Performance with Domain Knowledge*. Berlin: Springer. doi.org/10.1007/978-3-319-28971-7

Snidaro, L.; Vaci, J.; García, J.; Martí, E. D.; Jousselme, A. L.; Bryan, K.; Bloisi, D. D.; and Nardi, D. 2015. A Framework for Dynamic Context Exploitation. Paper presented at the 18th International Conference on Information Fusion. Washington DC, July 6–9.

Snidaro, L.; Visentini, I.; and Bryan, K. 2015. Fusing Uncertain Knowledge and Evidence for Maritime Situational Awareness via Markov Logic Networks. *Information Fusion* 21: 159–72. doi.org/10.1016/j.inffus.2013.03.004

Steinberg, N., and Bowman, C. L. 2009. Revisions to the JDL Data Fusion Model. In *Handbook of Multisensor Data Fusion*, edited by M. E. Liggins, D. L. Hall, and J. Llinas, 45–68. London: CRC Press.

Steinberg, N.; Bowman, C. L.; and White, F. E. 1999. Revisions to the JDL Data Fusion Model. In *Proceedings 3719: Sensor Fusion: Architectures, Algorithms, and Applications III*. Bellingham, WA: SPIE. doi.org/10.1117/12.341367

Strang, T., and Linnhoff-Popien, C. 2004. A Context Modeling Survey. Paper presented at the 1st International Workshop on Advanced Context Modeling, Reasoning and Management. Nottingham, UK, September 7–10.

Sycara, K.; Glinton, R.; Yu, B.; Giampapa, J.; Owens, S.; Lewis, M.; and C. Grindle. 2009. An Integrated Approach to High-Level Information Fusion. *Information Fusion* 10(1): 25–50. doi.org/10.1016/j.inffus.2007.04.001

Voss, P. 2007. Essentials of General Intelligence: The Direct Path to Artificial General Intelligence. In *Artificial General Intelligence. Cognitive Technologies, edited by* B. Goertzel and C. Pennachin, 131–57. Berlin: Springer. doi.org/10.1007/978-3-540-68677-4_4

White, F. E. 1987. Data Fusion Lexicon, Technical Report, San Diego, CA: Naval Ocean Systems Center.

Zadeh, L. A. 1983. The Role of Fuzzy Logic in the Management of Uncertainty in Expert Systems. *Fuzzy Sets and Systems* 11(1–3): 199–227. doi.org/10.1016/S0165-0114(83)80081-5

Lauro Snidaro received his PhD (2006) in computer science from the University of Udine, where he is now an associate professor. His main interests include data/information fusion, computer vision, machine learning and multimedia. Appointed Italian member of several NATO Research Task Groups on information fusion since 2003, Snidaro was the lead editor of *Context-Enhanced Information Fusion — Boosting Real-World Performance with Domain Knowledge,* published by Springer in 2016. He is editor of the fusion for signal/image processing and understanding area for the Elsevier *Information Fusion* journal, and of the high level fusion area for the ISIF *Journal of Advances in Information Fusion.*

Jesus Garcia is an associate professor in the Computer Science Department at the Universidad Carlos III de Madrid. His main research interests are computational intelligence, sensor and information fusion, machine vision, traffic management systems and autonomous vehicles. Within these areas, including theoretical and applied aspects, he has coauthored more than 10 book chapters, 60 journal papers and 180 conference papers. He has served on several advisory and programming committees in organizations for IEEE, ISIF, and NATO. He is the chair of the Spanish IEEE Chapter on Aerospace and Electronic Systems since 2013 and appointed Spanish member of several NATO-STO Research Groups.

James Llinas is a research professor emeritus and rounding director of the Center for Multisource Info Fusion at the State University of New York at Buffalo, New York, USA. He is an internationally-recognized expert in sensor, data, and information fusion, coauthored the first integrated book *Multi-Sensor Data Fusion,* and has taught and lectured internationally for more than 20 years on this topic. He has also coedited and coauthored the *Handbook of Data Fusion, Distributed Fusion for Net-Centric Operations,* and most recently, *Context Enhanced Information Fusion*.

Erik Blasch is a program officer at the Air Force Office of Scientific Research. He received a BS in mechanical engineering from Massachusetts Institute of Technology; MS in mechanical, health science, and industrial engineering from the Georgia Institute of Technology; MS in electronics, economics, and business from Wright State University; and a PhD in electrical engineering from Wright State University. He has compiled more than 750 papers, 21 patents, and 5 books focusing on robotics, information fusion, and man-machine systems, starting with the 1994 AAAI Robot Competition Winner (*AI Magazine* 16(2), 1995). He is an ISIF member, AIAA Associate Fellow, SPIE Fellow, and IEEE Fellow.