RC25254 (W1112-082) December 29, 2011
Computer Science

IBM Research Report

Quality of Information in Sensor Networks

Chatschik Bisdikian
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598 USA

Lance M. Kaplan
U.S. Army Research Laboratory
Adelphi, MD USA

Mani B. Srivastava
University of California
Los Angeles, CA USA

Research Division
Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

LIMITED DISTRIBUTION NOTICE: This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g., payment of royalties). Copies may be requested from IBM T. J. Watson Research Center , P. O. Box 218, Yorktown Heights, NY 10598 USA (email: reports@us.ibm.com). Some reports are available on the internet at http://domino.watson.ibm.com/library/CyberDig.nsf/home .

Quality of Information in Sensor Networks

CHATSCHIK BISDIKIAN, IBM Research
LANCE M. KAPLAN, U.S. Army Research Laboratory
MANI B. SRIVASTAVA, Univ. of California, Los Angeles

The increasing deliberate and/or ad-hoc deployment of sensor networks and the premeditated and/or opportunistic use of sensor-derived information provides enhanced visibility to everyday activities and processes that enables fast-paced decision making in personal, social, civilian, military, and business contexts. The effectiveness of the decisions made depends on the quality of the information gathered. In this paper, we highlight and build upon our recent work in the area of *quality of information* (QoI) for sensor networks. We present a quality-value layered definition of QoI, where the former relates to context-independent aspects and the latter to context-dependent aspects of an information product. Then we present a taxonomy of pertinent quality and value attributes anchored around a simple ontological relationship between the two. We, then, introduce a framework for assessing information products across their various attributes and ranking them based on the value they bring to a task. We close with a summary of the paper and a brief discussion of future directions for QoI research.

Categories and Subject Descriptors: C.2.2 [**Computer-Communication Networks**]: Network Protocols

General Terms: Design, Algorithms, Performance

Additional Key Words and Phrases: Wireless sensor networks, media access control, multi-channel, radio interference, time synchronization

ACM Reference Format:
Bisdikian, C., Lance, K. M., and Srivastava, M. B. Quality of Information in Sensor Networks ACM Trans. Sensor Netw. 9, 4, Article 39 (March 2010), 23 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION

We gather information to gain knowledge and build an understanding of things that interest us. For example, we gather information about a painter's life and the cultural and political situation surrounding it to gain a greater understanding and a deeper appreciation of the situations and events that influenced his or her artwork, or we collect information about a company's financial standing and actions to build an understanding of its future prospects. Quite often this knowledge and understanding aids us in some form of decision making and action taking, such as investing on the company

Author's addresses: C. Bisdikian is with IBM's T. J. Watson Research Center, Hawthorne, NY, USA, e-mail: *bisdik@us.ibm.com*; L. M. Kaplan is with the U.S. Army Research Laboratory, Adelphi, MD, USA, e-mail, *lkaplan@ieee.org*; M. B. Srivastava is with the Univ. of California, Los Angeles, CA, USA, *mbs@ee.ucla.edu*.

Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

(or the painting) or divesting from it. The degree to which we understand the world and the situations of interest depends on the *pertinence* of the information that we have access to, its accuracy, timeliness, completeness, provenance, and so on. In some broad sense, our understanding of situations depends on the *qualities* of the available information regarding the situation at hand.

We use sensor networks, such as wireless sensor networks (WSNs), to collect the pertinent (hopefully) information. The sensor-derived information (or simply, sensor information) will feed end-users such as humans, like an information analyst or a decision maker, or actuators, i.e., computer-controlled processes that act in response of the information they receive. Sensor information may pertain to the temperature distribution in a building whose environmental conditions need to be controlled, the vital signs and location of a patient whose health needs to be remotely monitored, the stress levels of a bridge whose structural health needs to be monitored, or the position, capabilities, and intentions of enemy troops, insurgents, etc., so that adequate preparations can be made. Sensor networks could be deployed purposely for supporting specific sensor-enabled applications or overlayed on top of existing "general-purpose" data networks as is the case with participatory sensing supporting crowdsourced-based applications.

The "goodness" of a network in performing its communication tasks is typically described by *quality of service* (QoS) attributes (bandwidth, delay, delay jitter, and loss probability) and the network's ability to attain levels for these attributes that are appropriate for a class of applications, such as best effort, constant or variable bit-rate. Analogously, one would expect that the goodness of a sensor network in performing its information-delivering tasks will be described through *quality of information* (QoI) attributes. Typically though quality for sensor networks has been thought as cross-layered extension of traditional QoS elements including routing topologies for efficient data distribution, energy efficiency and network lifetime, deployment coverage, data aggregation and in-network-processing, power control and bandwidth maximization, etc. All these are very important aspects of a successfully operating sensor network and will ultimately relate to its effectiveness as an information source supplying applications with desired information. However, these introvert, network-oriented quality aspects of sensor networks cannot capture in their entirety the extrovert, information-producing, application-oriented quality aspects of the network. For example, in *Internet of Things* [Gershenfeld et al. 2004] or military related applications rely heavily on highly heterogenous ad-hoc or planed sensory infrastructures. The primary information fusion challenge in these applications is the extraction of relevant information from the overwhelmingly large pools of data to enable a decision maker to determine the best course of action in light of the mission objectives. In other words, the challenge is to best transform *data to decisions* (D2D) [Blasch et al. 2011]. While several of these information producing sensor networks are composed of systems engineered for a specific purpose, e.g., object tracking systems in military networks, the technologies used in these systems are usually generic, and these systems could be modified as needed for other applications. Because not all decisions can be anticipated beforehand, it is desirable that the sensor systems and the networks that transport and transform the raw data to information are flexible enough to accommodate multiple applications in order to present the best information possible to the decision maker in context of mission objectives and operational constraints. To this end, QoI provides the underlying framework to enable methods to best discover, collect and fuse data sources in light of the information needs of the decision maker.

Looking into the extrovert quality aspects of sensor networks, this paper is organized as follows: In Section 2, we present our sensor network high-level model and terminology that serves as a backdrop for the rest of the paper. In Section 3, we "build" our quality-value layered definition of QoI and in Section 4 we present a collection of related quality and value attributes organized according to this definition. In Section 5, we introduce an information valuation framework based on the value attributes of a piece of information. We conclude in Section 6 with summary of the paper, some observations, and a brief discussion of possible research directions regarding QoI. Sections include background and related work as necessary both to set the stage and contrast our research pursues from past efforts.

2. THE SYSTEM MODEL

The fusion model proposed by the Joint Directors of Laboratories (JDL) Data Fusion Working Group comprises a hierarchy of fusion processes spanning from low level fusion of features that are used for object detection, classification, identification, and tracking, to higher levels of fusion for current situation awareness as well as future impact assessments [Steinberg et al. 1999; Llinas et al. 2004]. Higher level fusion requires a broad range of information that is disparate over time, space, and spectrum. This information is derived from lower level fusion of the source data. These different levels of information fusion and assessments could happen at different points in the end-to-end path between the information sources and the end-users, the analysts and decision makers. Nonetheless, lacking broader context, lower level fusion will typically occur closer to the sources (e.g., sensor nodes), while higher layer fusion will occur further away from these sources, either in transit within the networks connecting multiple sensing systems (sensors or sensor networks) or closer to the end-users where information from many different disparate sources can coalesce to provide the broadest possible coverage.

Along with the layered information processing fusion architecture afforded by the JDL model, in Figure 1 we take a systems architecture view of fusion; for most applications, the JDL model fits within the processing occurring within network nodes or at a central processor near the user. Specifically, Figure 1 presents two views of a sensor network deployment: (a) an *information flow* view is illustrated on the left, and (b) a *layered operations* view that is executed on these flows is illustrated on the right. Both views are applicable to an end-to-end system comprising sensors and sensor data transport/fusion/application system. On the layered stack, the data acquisition and application(s)/middleware layer are part of any end-to-end system. On the other hand, the data transport and system level fusion layers are optional layers that might not exist in highly integrated sensor-enabled systems, where the sensing system is integrated with its own application (e.g., an autonomous robot). In this case, the entire stack collapses down to two layers, the data acquisition and the application layers-where applications themselves may be responsible for any sensor data manipulation such as fusion. However, for highly distributed, multi-sensor systems feeding information to multiple applications, such as the aforementioned sensing tasks, every layer in the stack shown could be present.

In this set-up, the term *measurements* applies to the raw (measurable) data collected by the sensors, such as acoustic measurements, temperature measurements, spectrum measurements. Nevertheless, the information leaving these data sources and transported through the sensor and other networks are referred to as *sensor observations*. A sensor observation leaving a sensor node may not always represent raw measurements but could be the outcome of a local sensor fusion (made using the sensor measurements), signifying, for example, the presence of an object as detected by a particular sensor at a particular time instant, or that an object's color is red. Another example includes image processing where the raw measurements are the light levels measured on a focal plane array. The processed observations could be some low level computer vision enhancement or a processed observation such as a target detected at pixel (x, y) represent a line of bearing representing (a, e) degrees in azimuth and el-

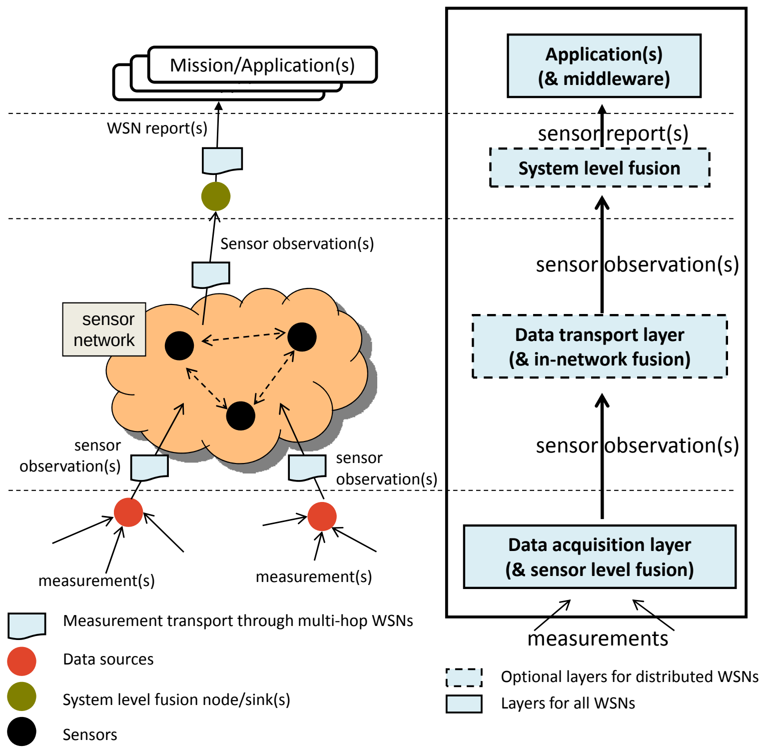

Fig. 1. Information flow and functional views of sensory information.

evation respectively. Note that the terminology used is in line with the definitions of measurements and observations in [Open Geospatial Consortium Inc. 2010], where an observation refers to the act of assigning a communicable descriptive representation (e.g., a number, term or other symbol) to a phenomenon, while a measurement relates to observations that pertain to measurable quantities (i.e., having an amount and a unit); observations related to categorical entities, such as color, are also referred to as "category observations" [Fowler 1996].

Sensor observations are processed at "system level" fusion sinks to produce system-level *sensor reports*, which are what applications experience coming (or reported) from the sensor network. For example, a sensor, through its local acoustic measurements, may observe the presence of a motorized four-wheel vehicle (a category-observation) in its vicinity, while the sensor network may confirm, at its level, and report to applications the presence of this vehicle and it may also report additional "objects" observed by other sensors in the network. It is entirely possible that, depending on the number of fusion layers between sensors and applications, observations and reports may represent semantically (and event syntactically) the same information content. For this reason, we would succinctly refer to these as sensor *information products*, with sensor nodes and networks producing information products (in the form of sensor observations) and applications consuming them (in the form of sensor-originated reports). An information product could represent a stream of data (e.g., temperature readings) from a particular source or a collection of sources, or individual event reports, e.g., location of an explosion, etc.

When a sensing system is designed and deployed in *closed-coupled* fashion with the corresponding application, the application designer has adequate (with some caveats, of course) knowledge and control of the operational context of the sensors, e.g., which sensors (or types of sensors) may be involved in an observation, what sensors are (or will be) trained to observe, when to observe, for how long to observe, how to express their observations syntactically, etc. In other words, the applications know what *their* sensors can provide and the sensing system knows what the needs of *their* applications are, i.e., the information products produced by the sensors and consumed by the application implicitly share a common playing field, or, what we refer to as, a *common context of operation* (CCO). However, such a naturally shared CCO cannot exist in *loosely-coupled*, open deployment environments where sensing systems and sensor-enabled applications could be designed, deployed, and operated at different times and by entirely different organizations populated by people of different expertise. Coping with unknown CCO could become a persistent reality when considering wireless and mobile senor networks where associations between sensors and applications can be ad-hoc and transient as well. Thus for the highly dynamic, open, late binding and rebinding cases, CCO needs to be established on-the-fly and information enrichment through QoI attributes and associated metadata can play a central role for this.

3. QUALITY OF INFORMATION

Robert M. Pirsing, in his philosophical novel "Zen and the Art of Motorcycle Maintenance: An Inquiry into Values" (William Morrow, 1974), said "...Even though quality cannot be defined, you know what it is ..." underscoring the easiness of understanding quality in general, i.e., having a gut feeling about it, but the difficulty in formally defining it. Such is the case with quality of information as well.

In the computerized world, information quality has been traditionally studied in the context of enterprise data management systems and processes, managing information stored in data warehouses to support business-level processes and decision making [Wang and Strong 1996]. The term has also been associated with the outcome of Web searches [Knight 2007], and with information fusion applications of detection, classification, identification, and tracking for military applications [Blasch et al. 2004; Blasch et al. 2010]. The most used definition of *information quality* (IQ) in these cases refers to information that is fit-to-use (for a purpose) –[Ehikioya 1999] considers a similar definition of processed data relevant to a user– and quality attributes such as accuracy, consistency, completeness, timeliness, are frequent terms in the lexicon of IQ. The terms "dimensions" and "characteristics" were also used in the literature instead of "attributes;" we will also use the term "metadata" later in the paper as well.

Both the traditional, business-centric and Web-search views of IQ apply to information that is consumable by human end-users, e.g., the employees of a company, its business strategy organization, a computer user submitting a query trough a Web-search engine, an intelligence analyst, see also [English 1999]. Information in these cases resides in structured (e.g., database) records or unstructured (e.g., text) documents, which are put together, retrieved, disseminated, and presented to humans at time-scales that reflect the human-centric processes. These can range from a few seconds for Web searches to days, weeks, and even more, when it comes to the collection of business-related information, e.g., the planing and execution of market analysis through customer interviews and the subsequent tabulation and analysis of the results. Furthermore, information in this case is considered to be an enterprise asset in that a business unit can plan and exercise control over the processes of collecting and normalizing information for storage and dissemination. This asset aspect of information has given rise to management processes for ensuring IQ such as the *to-*

tal data quality management (TDQM) program at MIT[1] and *total information quality management* (TIQM) system [English 2009], as well as to industry efforts to establish good data quality management practices for inter- and intra-business processes [West 2008].

The view that information is an enterprise asset, or more generally an organizational asset, that the user (or someone in the user's organization) can exercise a reasonable amount of control to improve its quality before it is delivered to a user is also implicitly shared by the JDL-related applications. The latter assume ownership of (or, at least, access privileges to) the sensory sources which allows the user to manage them to produce information of quality that satisfies the user's needs. Such applications are examples of the tightly-coupled systems we mentioned at the end of Section 2 that require careful planning before deployed in the field. However, the view of information as an organization asset does not account for cases where information sources are beyond the users control, such as the case with Web searches. Trust of a sensory source may not necessarily mean trusting only its sensing capabilities [Blasch et al. 2004], but also trusting the supplier of the information who also possibly owns the information-generating sensory source. In particular, with the emergence of loosely-coupled, dynamic application paradigms such as the *Internet of Things* (IoT) [Gershenfeld et al. 2004], crowd-sensing [Burke et al. 2006], mobile ad-hoc wireless networks, etc., gives rise to multi-modal, highly-distributed, multi-administrative, on-demand, ad-hoc, etc., sensor-and-actuator loosely-coupled application paradigms. These also include fast-paced *machine-to-machine* (M2M) applications that challenge the time-scales, and operational assumptions that information quality has been traditionally viewed under. Furthermore, while information may still be used to drive decision making, its end users may not always involve humans directly, or the decision making may not be a human-controlled business process.

Thus, we feel the need to broaden the scope of IQ to better reflect the much broader range of alternatives under which information can assume form, collected, stored, shared, used, etc. This broader scope does not intend to replace the traditional views of IQ, but rather augment and complement them. Our point of departure for this broader scope is the observation that a "fitness for use" view for information quality does not necessarily lend itself to multiple uses of information within an organization [English 2009] and that quality may exhibit *inherent* and *pragmatic* qualities, i.e., accuracy and value, respectively, to a business [English 1999].

The adherence to enterprise and business processes in the aforementioned IQ treatises is too restrictive in view of the emerging dynamic application paradigms. Instead, we prefer to consider information first on its own right and then how it relates to a process or an application context in general. Thus, we still maintain the notion that the application context within which information is to be interpreted and used plays a key role in assessing information quality. However, we recognize that the application space where a piece of information can be used can be very broad and possibly unknown at the time of its acquisition. For example, a particular image of a given resolution may, at present time, be useful for some applications and not useful for others, but this situation may reverse itself at some future time. Thus, we feel that a definition of quality of information (QoI)[2] needs to explicitly acknowledge this possibility. In particular, while attributes such as "timeliness," "completeness," "relevancy," etc., will certainly relate to the quality of an information product, we prefer to consider them as attributes that are

[1] See, http://mitiq.mit.edu/.

[2] In the sequel, QoI, instead of IQ, will be our preferred abbreviation as it follows the established "Qo[X]" template used in networked-related qualities, such as QoS and QoE, for quality of service and experience, respectively, see also [Stankiewicz et al. 2011].

Table I. Examples of quality and value related statements and questions.

Quality statements (quality facts)	**Quality-related questions** (value judgments)
There are 10% of measurements missing.	Is 3%, 25ft, 15 min good enough for my needs?
	<*accuracy* and *timeliness*>
The sensor measurement(s) has a 3% margin of error.	Who provides the location information?
	<*trustworthiness*>
The location information is accurate to within 25ft.	Does the sensor information covers my needs, e.g., content-wise, spatiotemporally?
	<*completeness*>
This (sensor-derived) information is 15 minutes old.	Do I need this high-resolution image?
	<*pertinence* or *relevancy*>
This is a 12 (4x3) megapixel image (high resolution).	Can I "consume" this piece of information, e.g., does it possess my desired syntactic/presentation form?
	<*readability*>

derived from the inherent quality properties of the product on a per use-case. Hence, a statement, for example, that an information product is not timely enough, should not penalize the product in perpetuity but rather only within the scope of a particular use. For other uses of the product, using again the inherent quality properties, alternative timeliness statements could be made.

More concretely, let us take a look at Table I–the table should be read a column at a time. On the left column, we have a collection of *quality statements* that can be ascribed to a piece of information, i.e., an information product, permanently as they describe innate properties of that product. On the right column, we have a collection of quality-related questions that answering them leads to *value judgments* regarding an information product. These value judgments, which are based on quality facts such as those on the first column, are ascribed to an information product on a case-by-case basis and only within the scope of the application that uses them and within the context of a particular use. Hence, a high-resolution image may be deemed relevant in one case and entirely irrelevant in another. Most importantly, when a judgment of "low value" is made for an information product, one should not necessarily be precluded from (or be biased) against ascribing a "high-value" rating to the same product ever again in the future.

The aforementioned split aspect of quality will carry through to our definition of QoI shortly. Before presenting the definition, though, we first present a few other quality definitions from industry standards bodies. ISO 9000 is a family of standards related to quality management systems. It defines quality (within its context) as:

— *ISO 9000*: Degree to which a set of inherent *characteristics* fulfills requirements.

ITU-T Rec. E.800 deals with *quality of service* in the provision of digital telecommunication services. It defines quality (within its context) as:

— *ITU-T Rec. E.800*: Totality of *characteristics* of a telecommunications service that bear on its ability to satisfy stated and implied needs of the user of the service.[3]

The IETF RFC 2386 deals with a framework for quality of service based routing in the Internet. A glossary entry in the RFC specifies quality (within its context) as:

— *IETF RFC 2386*: A set of service *requirements* to be met by the network while transporting a flow.

[3]This definition is from the 5th edition of Rec. E.800 [International Telecommunication Union 2008]. The following definition from the previous edition, published in 1994, appears frequently in the literature as well: "The collective effect of *service performances*, which determine the degree of satisfaction of a user of the service."

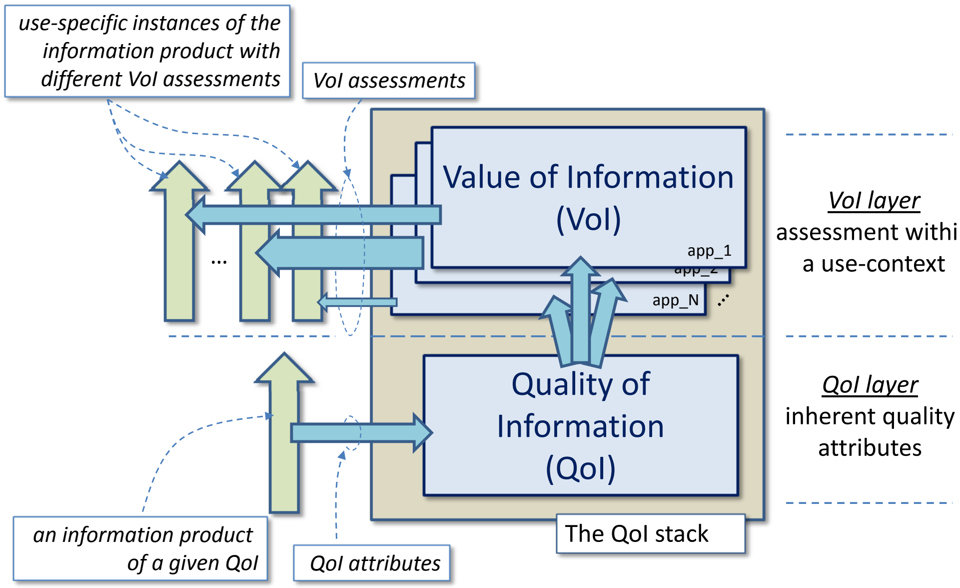

Fig. 2. The layering of QoI and VoI.

In the above statements, we have highlighted the terms "characteristics" and "requirements" as they point to the use of *dimensions* along which judgments of quality (within their respective contexts) are to be made. We adopt the same principle in our QoI definition. However, since we would like to be able to reconsider an information product as often as necessary, we do not ascribe a use-context to an information product from the outset (as the above definitions implicitly do). Instead, borrowing from the previous split view of quality, we explicitly separate the *ability of judging* an information product from the *outcome of judging*. The ability of judging relates to quality and quality attributes (the inherent characteristics) of an information product that feed the judging process. The outcome of judging relates to the value that a piece of information brings to an end-use. Using the quality of the available information, value judgments are made about this information on case-by-case basis.

One could envision quality and value represented by layers of an abstract stack with the value layer stacked on top of the quality layer. This is shown in Figure 2, where an information product, e.g., about the location of a truck, is annotated with QoI attributes, such as location error, time the location refers to, source owner/operator, etc. This product may be used by a number of applications, such as a surveillance application, a mobility modeling tool, or an intelligence analysis application. The information product with the given QoI may be valued differently in different use cases; this is indicated by the varying widths of the VoI assessment arrows in the figure.

Figure 3 adds a temporal aspect to the QoI/VoI stack of Figure 2. Specifically, the figure marks the time t_v at which an information product is to be valued within a

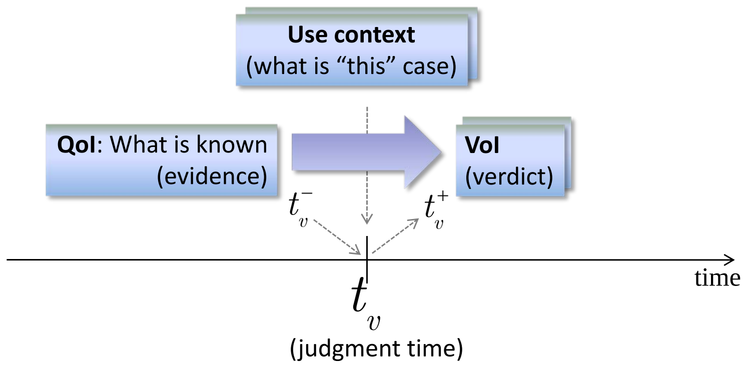

Fig. 3. Temporal relationship between QoI, use context, and VoI.

particular use context. All meta-information collected prior to that time (noted as t_v^-) that can feed into the valuation process represents the product's QoI. The QoI and the specifics of the use case results in the product's VoI at time t_v^+ (for this particular case). The figure also draws parallels with a trial process where a verdict is reached by analyzing the evidence available at judgement time within the context of a particular trial case; note that a piece of evidence may have different value at different trial cases.

With the above in mind, we define:

— *Quality of information* (QoI) the body of tangible evidence available (i.e., the innate information properties) that can be used to make judgments about the fitness-of-use and utility of information products.
— *Value of information* (VoI) is an assessment of the utility of an information product when used in a specific usage context.

Each definition relates to the corresponding layer in Figure 2, and the entire 2-layer stack represents the scope of the QoI domain. We acknowledge the abusive reuse of the term QoI that serves as the acronym for both the entire focus area of this paper as well as for the lower layer of the stack in the figure. Nonetheless, having the above definitions provides proper context when, say, an intelligence analyst declares that information products from a particular source are of low quality (or, low value) for her job. Incidentally, contrast her declaration with that of a digital camera evaluator declaring that the "superFancyCamera X11" (the source) takes high-quality pictures (the product), a declaration that is devoid of the use context for these pictures; this information is of high-value to those seeking to purchase a camera. Looking at the pictures, the analyst may too marvel the vibrant colors of the breathtaking sunset scenery captured by the camera. However, she may still ascribe a low value rating to them if they are taken at an unknown or irrelevant place and time for her job.

In addition to the layering of QoI and VoI in figure 2, [Thornley et al. 2009a] has provided a further refinement of VoI when considering "hybrid" end-to-end use cases, such as intelligence gathering, analysis and action taking and military missions. In these cases, information gathering involves a combination of both fast-paced (hard) and slower-paced (soft) sources, e.g., sensors and humans, feeding information to a *human-in-the-loop* decision maker, e.g., the commander of a mission. According to this refinement, value (hence, VoI) is perceived by the human decision maker, while the mission experiences the *utility of information* (UoI). Specifically, VoI reflects the confidence the decision maker could gain in making a decision by due time, when a particular information product becomes available, given the decision maker's cognitive state, including her current awareness of the situation, the effects of information overload, work overload, training and past experiences, etc. UoI is the impact that the information content of the product would have on the outcome of the mission should it be utilized. In the absence of the human unpredictability, VoI and UoI are expected to track each other, however, this may not always be true. For example, given his cognitive state, the decision maker may decide to act upon the availability of a piece of information that appears to reinforce a prior, biased belief of his and, hence, subjectively (and possibly unconsciously) perceive it as high value. However, the actions taken may have less than desired outcome for the mission, in which case, the seemingly high-valued piece of information for the decision maker had little, or even negative, utility for the mission. This refined view would be applicable to broader human-to-computer interactive systems, but our focus is on the automated and computerized portions of sensor-enabled applications and services and we will consider VoI and UoI indistinguishably.

We close this section by noting that sensor observations are used, typically, to populate the state of a model (or, abstraction) of the world of interest, e.g., provide for the location and velocity of a vehicle crossing an area under surveillance. In this case, *information* represents an estimate of that state; QoI represents the estimate's goodness measured along quality dimensions such as *accuracy* and *latency*, and influenced by the various processes that produced it, i.e., its *provenance*; and VoI represents its importance (with the given QoI) in a specific use-case.

In the next section we will delve further on the description of QoI and VoI though quality and value attributes.

4. THE 5WH PRINCIPLE AND QOI/VOI ATTRIBUTE ORGANIZATION

QoI represents the body of evidence available about an information product that can be used to make judgment statements about its value for a particular purpose. In this respect, the information content itself, e.g., a temperature measurement, or the scenery in a photograph, is not the only thing of importance. For example, Table II shows an information product regarding a temperature measurement whose contextual description becomes richer and richer as information about the information, i.e., the meta-information (or, *metadata*) representing the body of evidence, is added at each row. The added information allows us passing a better judgment regarding the (application-specific) value that the temperature measurement of 26°C could have. For example, after the sixth row in the table has been revealed, we can conclude that this measurement is of high-value for a historical data analysis of sorts (e.g., of one's office location), but of quite low-value for a real-time monitoring application. This value dis-

Table II. Metadata enriching a temperature measurement.

(1)	26°C											
(2)	26°C	(±1°C)										
(3)	26°C	(±1°C)	3^{rd}	floor								
(4)	26°C	(±1°C)	3^{rd}	floor	South	side						
(5)	26°C	(±1°C)	3^{rd}	floor	South	side	19	Skyline Dr.				
(6)	26°C	(±1°C)	3^{rd}	floor	South	side	19	Skyline Dr.	July	15	2011
(7)	[**26**]°C	(±[**1**]°C)	[$\mathbf{3^{rd}}$]	[**floor**]	[**South**]	[**side**]	[**19**]	[**Skyline Dr.**]	[**July**]	[**15**]	[**2011**]

tinction could not have been made with the body of evidence available up and including the fifth row.

We have repeated the last row of the table and bracketed its column entries to underscore another very important QoI attribute that of *provenance*. Provenance describes how the information at hand, including any related metadata, came to be, for example, the entries in the last row may have been randomly selected from a set of alternatives. Associated with provenance are also qualities of the source of information and the processes and entities that interacted with it. For example, source qualities may pertain to when was a sensor last calibrated, and by whom, or what random process was used to select the entries in the seventh row of Table II and how reliable is that process. Provenance thus relates, among others, to trust and reputation in the information source.

The added body of evidence is captured by a collection metadata ascribed to an information product instantiating various information attributes. These metadata may be persistently attached to the information product and "travel" together with it, or may be stored at metadata repositories and accessed on a per need basis. Since, seeking for a desired information product involves processing of its metadata, the choice between persistently attaching to or maintaining the metadata separately from their corresponding information products (or providing some in-between hybrid combinations) is a design trade-off between storage, bandwidth and responsiveness constraints.

In [Bisdikian et al. 2009a], we introduced the `5WH` primitive principle (`what`, `where`, `when`, `why`, `who`, and `how`) to systematically account for (meta-)information pertinent to an information product. The `what` primitive represents everything directly tied to the information content, like the measurement value, its units, its accuracy, etc., while the `where` and `when` primitives represent its physical context, i.e., the location and time that the information content relates to. These three primitives, which relate to the properties of the information content, are in line with the Semantic Sensor Web's (SSW) effort to enrich sensor data with spatial, temporal, and thematic annotations [Sheth et al. 2008]. In addition to the content-related primitives, the `who` and `how` primitives relate to the source of information. Specifically, the `who` primitive represents the information provider, like the provider's name, license ID (if one is required), trust level, etc., while the `how` primitive represents the specific sources and processes (e.g., the sensors and fusion algorithms) used to produce the information. In other words, the `who` and `how` primitives relate to the information provenance. Finally, the `why` primitive represents the "mission" for which an information product is needed, e.g., for securing a perimeter, and it forms the basis for VoI, i.e., of valuing an information product within the context of a specific use. Figure 4 summarizes the relationship between the `5WH` and the QoI/VoI stack.

In [Bisdikian et al. 2009b], we presented a first attempt to categorize QoI and VoI attributes and metadata inspired by the `5WH` principle. In example cases considered in [Bisdikian et al. 2009b], we presented QoI and VoI metadata classes expressed in UML data model abstractions, hence, naturally leaning towards organizing the attributes according to possible structural relationships when instantiating their cor-

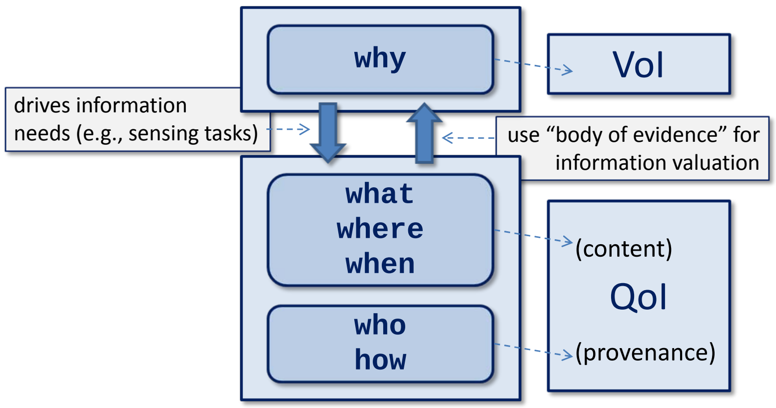

Fig. 4. The 5WH primitive principle for information attributes.

responding metadata. Despite the layered definition of quality and value, our initial organization had a few inconsistencies in naming, e.g., using `timeliness` (which is clearly a use-dependent judgment call) to describe a quality attribute such a `latency`, or considering `reputation` as a value attribute.

More recently, [Rogova and Bosse 2010] considered information quality in interactive computer-aided human decision-making systems and introduced an information quality ontology centered around the quality of the (a) information source, (b) information content, and (c) information presentation. While [Rogova and Bosse 2010] did not make the distinction between QoI and VoI, its concise ontology of pertinent quality attributes lent a helpful hand in our effort to refine more consistently our organization of QoI and VoI attributes that is not tied to the structural relationships in instantiations of the associated metadata. We note that the number of quality attributes considered has been quite generous at times–over 170 were noted at one point in [Wang and Strong 1996], before they were trimmed down to 20 representative quality dimensions (representing collections of similar attributes) which they were further grouped into 4 broad categories (accuracy, relevancy, representation, and accessibility). Furthermore, [Knight 2007] noted of 20 different information quality frameworks each with its own collection of attributes, dimensions, and categorizations. Hence, the conciseness of the ontology of quality attributes in [Rogova and Bosse 2010] was indeed quite attractive.

In refining our attribute organization, we started by abstracting the original metadata classes in [Bisdikian et al. 2009b] and then placed them in thematically related categories of attributes by borrowing nomenclature from the ontology in [Rogova and Bosse 2010] whenever appropriate. We then asked ourselves whether a particular category of attributes can be interpreted likewise independently of a use or not. We assigned the former categories to the QoI attribute "bucket" and the latter to the VoI bucket. We have considered only a subset of the categories in [Rogova and Bosse 2010] as the remaining ones were not relevant to our case or were more specialized versions of the categories considered, in which case they implicitly have also inherited the QoI vs. VoI categorization designations of their progenitors. Figure 5 shows the top ontological relationship of the refined attribute bucket categorization, where an information product is qualified by (see *has* in the figure) a collection of QoI attributes (and associated metadata) which together with an end-use context determine (*computes*) its value described by the VoI attributes (and associated metadata) assigned (*ascribed*) to the product for the particular use.

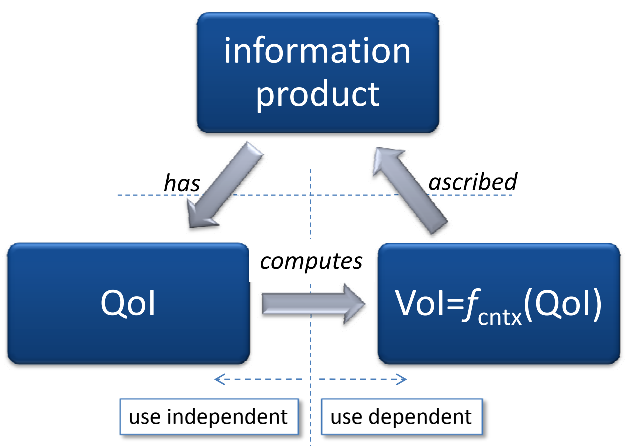

Fig. 5. The root information product/QoI/VoI ontological relationship.

Figure 6 shows the refined organization of QoI attributes in a taxonomy of QoI attribute categories. Having no direct counterpart in [Bisdikian et al. 2009b], the two-tone attribute categories of `content` and `accessibility` have been directly borrowed from [Rogova and Bosse 2010] and used as appropriate in our refined organization; also, the yellow-lettered categories of `accuracy`, `latency`, and `provenance` reflect the

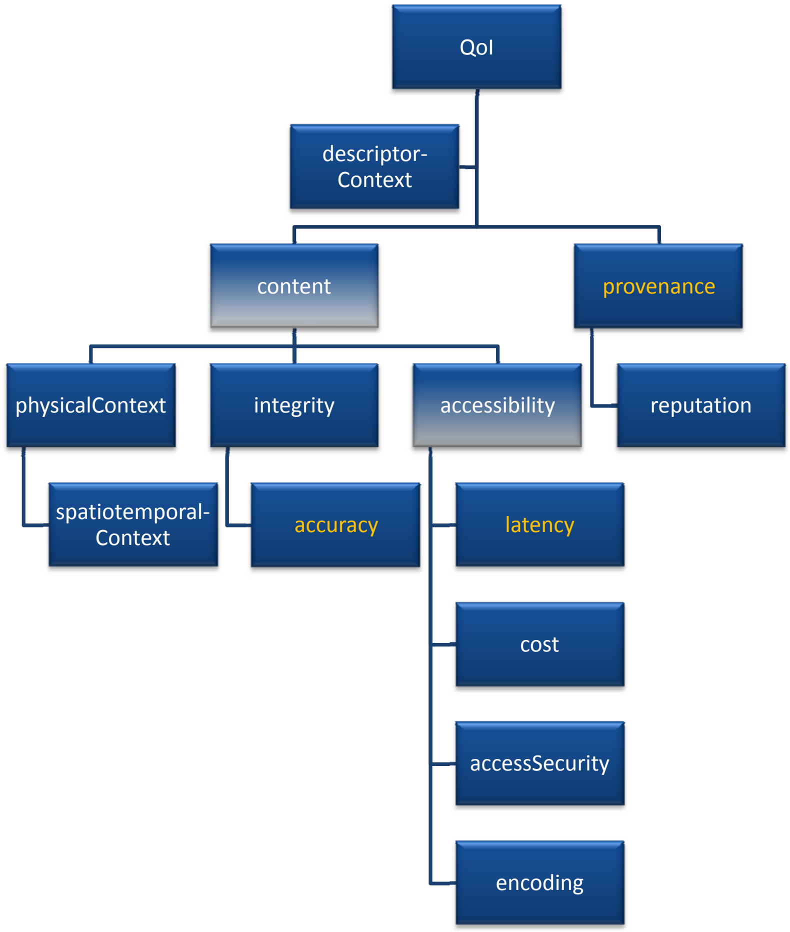

Fig. 6. The QoI attribute taxonomy.

closing comments in Section 3. First of all, the attributes and associated metadata[4] are interpreted within a `descriptorContext` that plays role similar to our original `AppDomainCntx` UML class for QoI in [Bisdikian et al. 2009b], and describes the syntactic and semantic rules that the rest of the metadata may follow, e.g., they are encoded according to SensorML [Open Geospatial Consortium Inc. 2007]. The rest of the QoI attributes are organized under two broad categories, those applicable to `content` and those to `provenance`. These two broad categories correspond directly to the entries in the QoI block in Figure 4.

The `content` category comprises three subcategories pertaining to `physicalContext`, `integrity`, and `accessibility`. The `physicalContext`, similar to our original `Context` in [Bisdikian et al. 2009b], includes the `spatiotemporalContext` of the information, i.e., the time and space horizon over which this information product pertains and is

[4]In the sequel, we will not distinguish between the terms attributes and (associated) metadata and will use them interchangeably.

valid for. The `integrity` relates to imperfections in the reported content values and has been extensively subdivided in [Rogova and Bosse 2010] under two main categories: `uncertainty` and `imprecision`. For our sensor-derived information, we particularly consider `accuracy`, contained under `imprecision` in [Rogova and Bosse 2010], expressing any knowledge about error behavior, such as error range, variance, etc. Thus, while we previously had entirely different (UML) classes for `integrity` and `accuracy` under a superclass for QoI attributes, since these are thematically related we have (re)categorized their relationship accordingly. The `accessibility` category encompasses attributes that relate to the accessing the information product, such as the `cost` for accessing it, in both monetary and non-monetary terms, e.g., energy consumption, as well as the time it takes to retrieve the information (`latency`) including the time for tasking sensors and sensing, when new sensor measurements are required, `accessSecurity`, and `encoding` describing how the information is represented or can be accessed, e.g., a Java serialized object; the latter is analogous to our original `QoIFormatAttr` in [Bisdikian et al. 2009b].

The `provenance` category relates to the entire end-to-end source-to-sink path that the sensor-derived information followed. Hence, `provenance` relates not only to information sources, as considered in [Rogova and Bosse 2010], but also to any processing, e.g., fusion, that happened from the moment the information was originally "sensed" until it arrived to its destination. While one could consider the output of information processing to represent a new source of information, we prefer considering `provenance` as capturing the entire (known) chain of information processing. Within `provenance`, we specifically note the `reputation` of the source (or process) that describes a *publicly* held opinion about the source's competence, truthfulness, etc., in providing content at the stated quality levels. We contrast this with `trust` (see in VoI later on) which results from direct or indirect interactions of information *users* with sources (or processes);[5] hence, we categorize `trust` under VoI. The `reputation` is one of the categories whose scores may evolve with time. Such categories may be populated with the latest available scores just prior to a valuation time, see time t_v^- in Figure 3.

Figure 7 shows the refined organization of VoI attributes in a taxonomy of VoI attribute categories. We explicitly mark the `end-useContext`, similar to our original `AppDomainCntx` class for VoI in [Bisdikian et al. 2009b], that forms the basis for the value judgments passed to the `content` and `provenance` of an information product. The `relevance` attribute category relates to how "close" (or, complete) the information content provided is to the one requested. It pertains to the physical and thematic *coverage* of the information product, the former represented by the `spatiotemporalRelevance` and driven by the `spatiotemporalContext` QoI attribute and the latter represented by the `thematicRelevance`[6] and driven by the original query that resulted in the specific information product. For example, consider a query about "images of public-use vehicles (buses, taxis) passing intersection $\mathcal{X}$ during morning and afternoon rush-hours," and an information product in response that only provides images of public-transportation buses (`thematicRelevance`) at intersection $\mathcal{Y}$, which precedes intersection $\mathcal{X}$, during morning rush-hours (`spatiotemporalRelevance`). Ref. [Tychogiorgos and Bisdikian 2011] considered the spatial relevancy based on the degree of overlapping that there is between the spatial characteristics of the desired and provided information. The `integrity` attribute category (including `accuracy`) relates to judgments made based on the corresponding QoI attribute, i.e., the degree to which the

[5]Indirection could result from trust transitivity, where one trusts (to a degree) somebody else's trust level for a source.
[6]This corresponds to the `completeness` typically found in the literature. We prefer using `thematicRelevance` instead as it describes what is about more clearly.

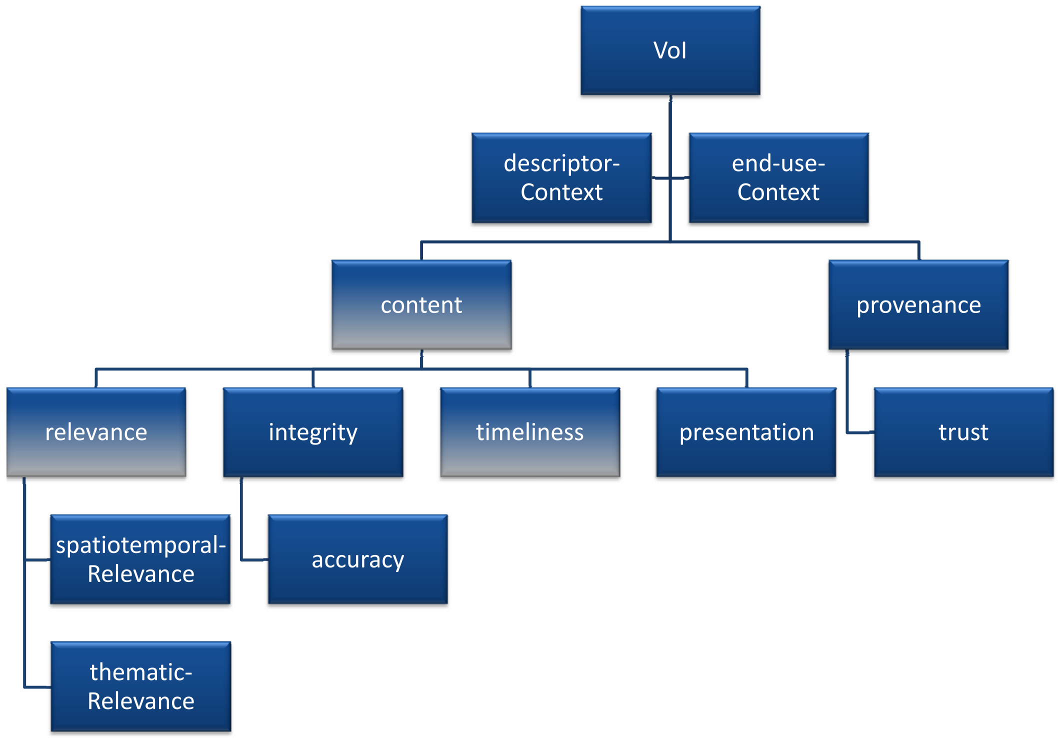

Fig. 7. The VoI attribute taxonomy.

integrity of the information product is satisfactory for the task at hand. Likewise, the timeliness relates to judgments made regarding the timely availability of the information product, hence, it closely relates to and based on the accessibility QoI attribute category. The presentation category, similar to our original VoIConvenienceAttr class in [Bisdikian et al. 2009b], pertains to how information is organized and presented to the user. Finally, under provenance, we specifically note trust which, as discussed earlier, is based on the disposition of information users against (or in favor of) the source of the information product that may result from past direct or indirect interactions with the source.

In closing, we note that in [Rogova and Bosse 2010] presentation was, for the purpose of that paper, at par with content and source (the provenance in our case). We could have placed presentation likewise in the VoI attribute taxonomy, or alternatively associate it with relevance and specifically thematicRelevance; in fact, in [Rogova and Bosse 2010] completeness (our thematicRelevance) was considered under presentation. However, we prefer to associate presentation directly with content, since it clearly relates to it and how it is *experienced* by end-uses. Also, doing so preserves the same two-prong high level content/provenance structure we have for the QoI attribute organization. However, we distinguish presentation from relevance which we consider relating only to the actual information content but not necessarily its representation.

5. VOI ASSESSMENT AND THE ANALYTIC HIERARCHY PROCESS

Assessing the value of an information products allows one to rate the fitness and utility of an information product to a task. More importantly, it allows determining which products are useful, ranking the useful products, and selecting those that can best support the information needs of users, especially in the presence of resource constraints, monetary or otherwise. The study of value of information has its roots in the theory of decision making. It has been related to the cost of acquiring an information product in order to reduce uncertainties due to insufficient knowledge during decision making and, also, the benefit derived from the ensuing decision and actions taken [Howard 1966; 1968]. Here, decision making will typically refer to business related actions (invest, divest, store raw material, sell in particular market segments, compete for a contract, initiate an ad campaign, and so on) and, hence, costs and benefits relate to moneys that are ultimately expected (probabilistically) to be gained or lost as a result of the decisions taken. More recently, [Hubbard 2007], considering again business actions, defined the *expected value of information* (EVI) as the difference in the *expected opportunity loss* (EOL)–the average loss in monetary benefit in the chance that the decision made was wrong–before the information product was available to the decision maker and after it became available. The implication here is that one should not pay more for acquire information than the reduction in EOL that it can bring.

The above cases deal with static situations where simple probabilities suffice to describe the related models. However, with evolving missions where cost and benefits change over time more, elaborate are more appropriate. Considering an intelligence, surveillance and reconnaissance (ISR) military mission, [Thornley et al. 2009a; Thornley et al. 2009b] considers combining stochastic models describing: (a) the mission physics, including the geography, sensing assets that are deployed, the sensing process, friend and foe mobility models, etc; (b) the intelligence service; (c) the situational awareness; (d) the decision makers; and (e) the actor. Applying performance evaluation process algebra (PEPA) techniques[7] a continuous time Markov chain (CTMC) abstract model of the system is built. The CTMC model together with simulations is used to estimate the possible mission outcomes along timeliness capturing different operational realities and initial conditions and capture the value that the availability of various information product might bring to the decision maker at different times.

The PEPA approach provides an assessment of VoI by stochastically recreating entire mission storylines off-line involving the aforementioned five models. Thus, it can provide fine-grained analysis of VoI tuned to the specifics of each different mission it models. For more timely assessment of VoI, especially when considering only the automated and computerized portions of sensor-enabled applications and services, techniques that consider assessment of VoI across its various attributes have been suggested. As our interest is in such automated systems, we draw inspiration from [Rogova and Bosse 2010] which provides an extensive discussion on "... assessing the values of information quality ... " suggesting the possible use of multi-criteria decision-making techniques for assigning a numerical value to an overall VoI measure.[8] Specifically, [Rogova and Bosse 2010] considers relative attribute weights to average assessment scores for the pertinent information attributes, where the weights are to be determined by experts in various application domains.

Weighted-averages along various attributes is a natural approach and has been used elsewhere too, for example in [Hossain et al. 2007] in estimating quality for a surveil-

[7] PEPA is a technique for modeling systems comprising concurrent processes that cooperate to achieve the behavior of the system [Gilmore and Hillston 1994].

[8] Ref. [Rogova and Bosse 2010] uses the term "quality measure" instead; however, "VoI measure" is more appropriate in our work.

Table III. VoI attribute relative scores and weights.

	relevance	integrity	timeliness	weight w
relevance	1	2	3	**0.5472**
integrity	1/2	1	1/2	0.1897
timeliness	1/3	2	1	0.2631

lance application along various quality factors. Likewise, we also consider weighting VoI attributes to evaluating an overall VoI score for an information product. By our definition of VoI, VoI attribute weights are use-specific, for example, `integrity` may weigh more than `timeliness` for a historical data analysis application, while for a real-time monitoring application be the other way around. To facilitate the process of developing weights, we introduce a repeatable assessment framework which can be used to systematically derive these weights by the experts in the various application domains.

5.1. A VoI valuation framework using AHP

Our VoI assessment framework is based upon the well-established multi-criteria decision-making technique called *analytic hierarchy process* (AHP) [Saaty 1990; Bodin et al. 2005]. AHP can be applied to produce attribute valuation weights systematically based on pair-wise comparisons of valuation criteria (the VoI attributes, in our case) and to ultimately rank information products. The name implies that it can be applied hierarchically, which for VoI means we can repeat the attribute valuation process at successive levels of the taxonomy tree in Figure 7 and in the end compose an overall VoI value. Next, we present an instance of the framework through a simple, multi-tier example.

Suppose we need to make VoI decisions based on information `content` and its three attribute subcategories `relevance`, `integrity`, and `timeliness`. We (i.e., the domain expert) start by populating Table III with valuation scores by comparing the attributes of interest in pairs and noting how the attribute in a row entry is "valued" relative to the attribute in a column entry for a particular use-case.[9] The scores are interpreted as follows: 1→*equal*; 3→*moderate*; 5→*strong*; 7→*very strong*; 9→*extreme*; in-between integer values are permitted and the inverse values (1/3, 1/5, etc.) imply reversing the comparison order. For example, based on the scores in the table `relevance` is somewhat more valued than `integrity` (a "2") and it is moderately more valued than `timeliness` (a "3").

The last column in Table III contains the weight vector w with the weights to be assigned to each of the attributes. It corresponds to the normalized principle eigenvector of the pairwise comparison matrix, i.e., the eigenvector of unit magnitude that corresponds to the largest eigenvalue. These weights can now be used to weigh the `relevance`, `integrity`, and `timeliness` scores that experts assign to specific information products–these scores are assumed to be normalized as well–hence, providing an overall VoI value for the products. We have marked with bold the highest ranked attribute, which, for this example, is the `relevance`.

The AHP-based framework can be used again for ranking products. For example, suppose that there are three information products, $\mathcal{A}$, $\mathcal{B}$, and $\mathcal{C}$, with pair-wise scores comparing how each product fares against the others for each of the VoI attributes of interest. These comparisons can be based on the metadata scores of the various QoI attributes for each of the information products, e.g., assess, within the particular use context, how much better the `relevance` of one product is relative to another us-

[9]AHP uses relative instead of absolute scores making it easier to accommodate both hard (objective/measurable) and soft (subjective) criteria.

Table IV. Relative scores of three products along the various VoI attributes and weights.

	`relevance`				`integrity`				`timeliness`			
	$\mathcal{A}$	$\mathcal{B}$	$\mathcal{C}$	weight w_1	$\mathcal{A}$	$\mathcal{B}$	$\mathcal{C}$	weight w_2	$\mathcal{A}$	$\mathcal{B}$	$\mathcal{C}$	weight w_3
$\mathcal{A}$	1	2	1/3	0.2809	1	2	3	**0.5396**	1	1/2	2	0.3313
$\mathcal{B}$	1/2	1	1	0.2552	1/2	1	2	0.2970	2	1	1/3	0.2894
$\mathcal{C}$	3	1	1	**0.4638**	1/3	1/2	1	0.1634	1/2	3	1	**0.3793**

Table V. Evaluation and rank of the information products.

	`relevance` (0.5472)	`integrity` (0.1897)	`timeliness` (0.2631)	*rank* (r)
$\mathcal{A}$	0.2809	0.5396	0.3313	0.3433
$\mathcal{B}$	0.2552	0.2970	0.2894	0.2722
$\mathcal{C}$	0.4638	0.1634	0.3793	**0.3846**

ing their QoI `spatiotemporalContext` attribute. As before, the same importance scale is used and weights based on the primary eigenvector are derived and summarized in Table IV. We mark with bold the highest ranked product for each of the three attributes. Finally, in Table V we derive the product rank r by properly combining the weighs, $r = [w_1|w_2|w_3] \cdot w$, which marks product $\mathcal{C}$ as the one providing the most value for the use-case at hand.

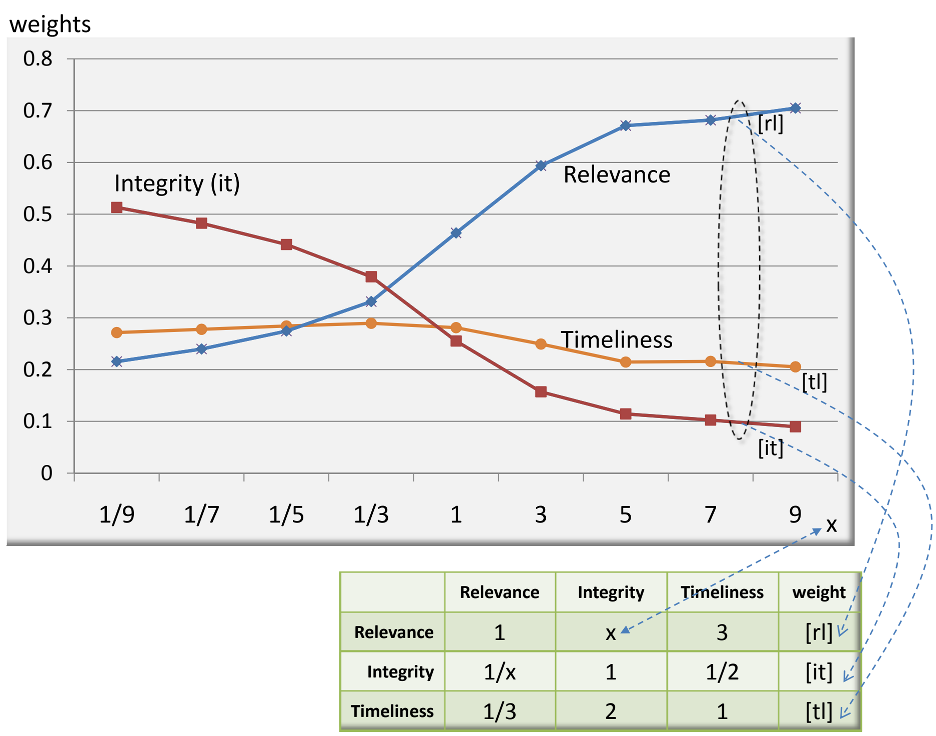

Fig. 8. VoI attribute weights for various `integrity` vs. `relevance` scores.

We have repeated the above process for various `integrity` vs. `relevance` scores while keeping the rest of the relationships as per Table III. The results are shown in Figure 8, where, as expected, `relevance` and `integrity` go opposite ways as the score between them changes, while `timeliness` remains relatively unaffected. Assuming the same product vs. VoI attribute scores as in Table IV, the corresponding valuations and ranking for the information products are shown in Figure 9. Since, `relevance` and `integrity` weigh significantly on information products $\mathcal{C}$ and $\mathcal{A}$, respectively, the valuations for these products mirror the behavior of these VoI attributes, while the valuation of product $\mathcal{B}$ remains again relatively unaffected.

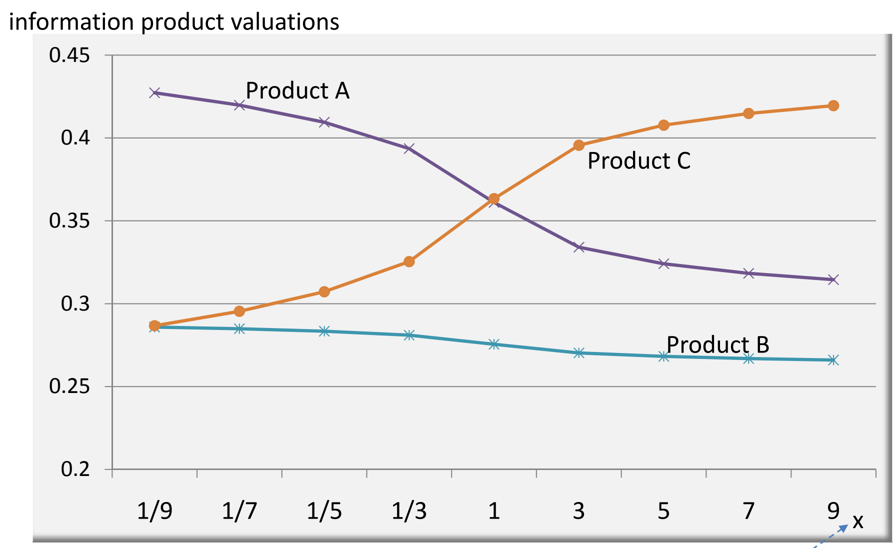

	Relevance	Integrity	Timeliness
A	0.2809	0.5396	0.3313
B	0.2552	0.2970	0.2894
C	0.4638	0.1634	0.3793

	Relevance	Integrity	Timeliness
Relevance	1	x	3
Integrity	1/x	1	1/2
Timeliness	1/3	2	1

Fig. 9. Evaluation of the information products for various `integrity` vs. `relevance` scores.

6. CONCLUDING REMARKS

The increasing availability of real-time and/or archived sensor-originated information is at the heart of the next information revolution where fast-paced decision making is supported by the enhanced visibility that up-to-the minute access to information from a diverse set of sensory sources provides. In an environment where applications may bind to information sources on demand the quality of the available information (QoI) plays an important role on the effectiveness of the decisions taken. Going beyond the operational aspects of sensor networks, such as QoS and coverage, QoI relates to the information characteristics of sensory information, such as accuracy, latency, and provenance. Building upon past industry definitions for quality, in this paper we presented a layered definition of quality and value of information, where the latter depends on the former to assess the potential utility that a particular information product brings to the task at hand. Then, refining our prior work on QoI metadata, we presented a taxonomy of QoI and VoI attributes grouping them under two major

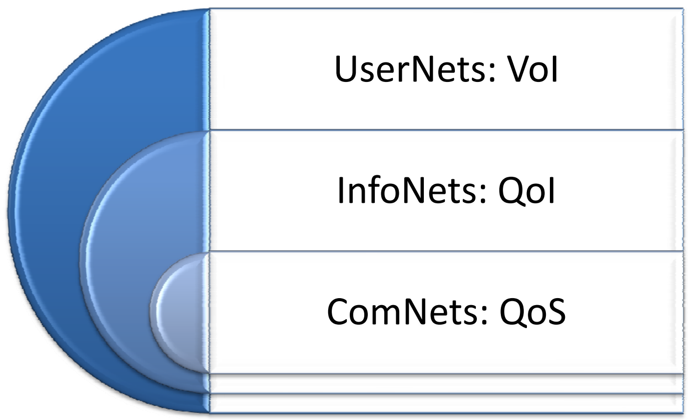

Fig. 10. The QoS/QoI/VoI relationship.

categories in each case, one for attributes related to the information content and one for attributes related to its provenance. We do not preclude further enrichment of the QoI and VoI taxonomies, however, we expect these to be in depth extensions maintaining the two-prong `content/provenance` structure at the highest level. Finally, we introduced a value assessment framework based on the analytic hierarchy process for multi-criteria decision making.

It is expected that, in general, information will be retrieved directly from the sensor networks or retrieved from sensory-data repositories (on the Web or elsewhere), then transported over a variety of networks (LANs, MANs, WANs, MANETs, public, private, you name it), and (possibly) further processed en route to the end-users. Hence, it is natural to expect that the QoS delivered by the communication networks could impact the quality of the information flows they carry, e.g., delay the information, reduce the amount of information delivered, and hence increase, say, the uncertainty in the inferences that can be made with it, which, in turn, may affect the value gained by the end-user consuming this information.

We summarize these points in Figure 10, where *UserNets* (the end-uses or end-users) depend on *InfoNets* (the information acquisition, storage and processing facilities) to provide the necessary information which, in turn, depend on *ComNets* (the information transport facilities) to deliver this information with some QoS characteristics, which impacts the information's QoI, which affects the VoI experienced by the end-users.

We close by briefly highlighting a few example directions that QoI-related research can pursue to further advance the state-of-the-art in the area. Firstly, there is a need to succinctly describe an end-user's information needs that can be easily communicated to sensory information providers on demand. Secondly, there is a need to investigate the relationship between the QoI produced by sensory-systems and the operational characteristics of these systems including the impact on QoI of faulty sensor operation, resource constraints, network performance, etc. Thirdly, there is a need to develop systematic approaches for evaluating the value to a sensor-dependent end-use that sensory-information of given quality could provide; these approaches should be easily reproducible and adaptable to the diverse (and unknown at the moment) variety of such end-uses. Fourthly, there will be the need to combine the above to develop effective management techniques for sensor networks that focus not only on improving traditional QoS but QoI and VoI for the applications they support. We believe that building upon the QoI/VoI thesis presented in this paper, and pursing additional research directions such as the ones highlighted above, see also the survey in [Sachidananda et al. 2010]–we have already been investigating aspects along these directions–will result in a critical and valuable body of art that would significantly advance the usability and ease of deployment and operation of sensor-based systems such as the ones envisioned in an *internet-of-things* world.

REFERENCES

BISDIKIAN, C., BRANCH, J., LEUNG, K. K., AND YOUNG, R. I. 2009a. A letter soup for the quality of information in sensor networks. In *IEEE Information Quality and Quality of Service (IQ2S'09) Workshop (in IEEE PerCom'09)*. Galveston, TX, USA.

BISDIKIAN, C., KAPLAN, L. M., SRIVASTAVA, M. B., THORNLEY, D. J., VERMA, D., AND YOUNG, R. I. 2009b. Building principles for a quality of information specification for sensor information. In *12th Int'l Conf. on Information Fusion (FUSION'09)*. Seattle, WA, USA.

BLASCH, E., VALIN, P., AND BOSSE, E. 2010. Measures of effectiveness for high-level fusion. In *13th Int'l Conf. on Information Fusion (FUSION'10)*. Edinburgh, UK.

BLASCH, E. P., PRIBILSKI, M., DAUGHTERY, B., ROSCOE, B., AND GUNSETT, J. 2004. Fusion metrics for dynamic situation analysis. In *Proc. SPIE, Signal Processing, Sensor Fusion, and Target Recognition XIII*, I. Kadar, Ed. Vol. 5429. Orlando, FL, USA.

BLASCH, E. P., RUSSELL, S., AND SEETHARAMAN, G. 2011. Joint data management for MOVINT data-to-decision making. In *14th Int'l Conf. on Information Fusion (FUSION'11)*. Chicago, IL, USA.

BODIN, L. D., GORDON, L. A., AND LOEB, M. P. 2005. Evaluating information security investments using the analytic hierarchy process. *Communications of the ACM 48,* 2, 79–83.

BURKE, J., ESTRIN, D., HANSEN, M., PARKER, A., RAMANATHAN, N., REDDY, S., AND SRIVASTAVA, M. B. 2006. Participatory sensing. In *World Sensor Web Workshop (in ACM Sensys'06)*. Boulder, CO, USA.

EHIKIOYA, S. A. 1999. A characterization of information quality using fuzzy logic. In *18th Int'l Conf. of the N. American Fuzzy Information Processing Soc. (NAFIPS)*, R. N. Dave and T. Sudkamp, Eds. New York, 635–639.

ENGLISH, L. P. 1999. *Improving Data Warehouse and Business Information Quality: Methods for Reducing Costs and Increasing Profits*. Wiley.

ENGLISH, L. P. 2009. *Information Quality Applied: Best Practices for Improving Business Information, Processes and Systems*. Wiley.

FOWLER, M. 1996. *Analysis Patterns: Reusable Object Models*. Addison-Wesley.

GERSHENFELD, N., KRIKORIAN, R., AND COHEN, D. 2004. The Internet of Things. *Scientific American*.

GILMORE, S. AND HILLSTON, J. 1994. The pepa workbench: A tool to support a process algebra-based approach to performance modelling. In *7th Int'l Conf. on Modelling Techniques and Tools for Computer Performance Evaluation*, G. Haring and G. Kotsis, Eds. Lecture Notes in Computer Science Series, vol. 794. Springer-Verlag, Vienna, Austria, 353–368.

HOSSAIN, M. A., ATREY, P. K., AND SADDIK, A. E. 2007. Modeling quality of information in multi-sensor surveillance systems. In *IEEE Data Engineering Workshop on Ambient Intelligence, Media, and Sensing (Part of IEEE ICDE 2007)*. Instabul, Turkey, 11–18.

HOWARD, R. A. 1966. Information value theory. *IEEE Trans. on Systems Science and Cybernetics SSC-2,* 1, 22–26.

HOWARD, R. A. 1968. The foundations of decision analysis. *IEEE Trans. on Systems Science and Cybernetics SSC-4,* 3, 211–219.

HUBBARD, D. W. 2007. *How to measure anything: Finding the value of "intangibles" in business*. Wiley.

International Telecommunication Union Sept., 2008. *Rec. ITU-T E.800: Quality of telecommunication services: concepts, models, objectives and dependability planning; Terms and definitions related to the quality of telecommunication services*. International Telecommunication Union.

KNIGHT, S.-A. 2007. User perceptions of information quality in world wide web information retrieval behavior. Ph.D. thesis, Edith Cowan University.

LLINAS, J., BOWMAN, C., ROGOVA, G., STEINBERG, A., WALTZ, E., AND WHITE, F. 2004. Revisiting the JDL data fusion model II. In *7th Int'l Conf. on Information Fusion (FUSION'04)*. Stockholm, Sweden.

Open Geospatial Consortium Inc. July 17, 2007. *OpenGIS Sensor Model Language (SensorML) Implementation Specification*. Open Geospatial Consortium Inc., (Ref. Num. Doc: OGC 07-000, ver. 1.0.0), M. Bots and A. Robin (ed.).

Open Geospatial Consortium Inc. Oct. 11, 2010. *Geographic Information: Observations and Measurements*. Open Geospatial Consortium Inc., (OpenGIS Proj. Doc: OGC 10-004r3 and ISO 19156), Simon Cox (ed.).

ROGOVA, G. L. AND BOSSE, E. 2010. Information quality in information fusion. In *13th Int'l Conf. on Information Fusion (FUSION'10)*. Edinburgh, UK.

SAATY, T. L. 1990. How to make a decision: The analytic hierarchy process. *European Journal of Operational Research 48,* 1, 9–26.

SACHIDANANDA, V., KHELIL, A., AND SURI, N. 2010. Quality of information in wireless sensor networks: A survey. In *15th Int'l Conf. on Information Quality (ICIQ 2010)*. Little Rock, AK, USA, 193–207.

SHETH, A., HENSON, C., AND SAHOO, S. S. 2008. Semantic Sensor Web. *IEEE Internet Computing*.

STANKIEWICZ, R., CHOLDA, P., AND JAJSZCZYK, A. 2011. QoX: What is it really? *IEEE Communications Magazine*.

STEINBERG, A. N., BOWMAN, C. L., AND WHITE, F. E. 1999. Revisions to the JDL data fusion model. In *Proc. of the SPIE, Sensor Fusion: Architectures, Algorithms, and Applications III*. Vol. 3719. Orlando, FL, USA.

THORNLEY, D. J., GILLIES, D. F., AND BISDIKIAN, C. 2009a. Toward mission-specific service utility estimation using analytic stochastic process models. In *Proc. SPIE, Modeling and Simulation for Military Operations IV*, D. A. Trevisani, Ed. Vol. 7348. Orlando, FL, USA.

THORNLEY, D. J., YOUNG, R. J., AND RICHARDSON, J. P. 2009b. Toward mission-specific service utility estimation using analytic stochastic process models. In *Proc. SPIE, Intelligent Sensing, Situation Management, Impact Assessment, and Cyber-Sensing*, J. F. Buford, G. Jakobson, S. Mott, and M. J. Mendenhall, Eds. Vol. 7352. Orlando, FL, USA.

TYCHOGIORGOS, G. AND BISDIKIAN, C. 2011. Selecting relevant sensor providers for meeting "your" quality information needs. In *12th Int'l Conf. on Mobile Data Management (MDM 2011)*. Lulea, Sweden.

WANG, R. Y. AND STRONG, D. M. 1996. Beyond accuracy: What data quality means to data consumers. *Journal of Management Information Systems*.

WEST, M. 2008. ISO 8000, standards for data and information. In *Data Management & Information Quality*. London, UK.