An Ontology of Information for Emergency Management

Antony Galton
University of Exeter, UK
apgalton@ex.ac.uk

Michael Worboys
University of Maine, USA
worboys@spatial.edu.maine

ABSTRACT

The next generation of information systems for emergency management will be based on information provided by large and diverse collections of sensors, including information supplied by human volunteers. Consequently there is more than ever a need to provide solutions to the integration question, so that the Common Operating Picture can truly and effectively provide the unified view required of it. This paper describes some work on the ontology of information that can contribute to a solution of the integration problem. To set the stage, the paper discusses the relevance of information integration to emergency management, and then goes on to describe a project that provided the catalyst for this work. Later sections introduce ontological research and proceed to use it to lay the foundations for an ontology of information. In the final sections we indicate how such an ontology can be used in the context of emergency management.

Keywords

Ontology, information

INTRODUCTION

The ultimate goal of any emergency management system is to improve *situation awareness*. This will result in improved perception of environmental elements; improved understanding of their meaning; enhanced ability to forecast the status of situational elements in surrounding areas and in future time; and understanding of how local decisions may impact goals. Improved situation awareness may be provided through the mechanism of the *common operating picture* (COP). The COP provides a single, unified display of relevant information as well as forming a framework in which collaborative planning can take place. In addition to providing a unified view of a situation, the COP should enable different perspectives on it, depending on the function of the viewer (e.g. whether a data collector, fire chief, or logistical manager in command of relief materials distribution).

The next generation of information systems for emergency management will be based on information provided by large and diverse collections of sensors. With the increased functionality and availability of communication systems, emergency management information systems will also be able to take advantage of information supplied by human volunteers. With this diversity of information sources and types, there is more than ever a need to provide solutions to the integration question, so that the COP can truly and effectively provide the unified view required of it.

This paper describes some work on the ontology of information that can contribute to a solution of the integration problem. Information is diverse in both its types and its generators, and both forms of diversity, along with their appropriate quality measures, need to be considered. Indeed, the question of what information is has no easy answer, but an answer needs to be attempted before the later issues can be addressed.

To set the stage, the paper motivates the research area by discussing the relevance of information integration to emergency management, and then goes on to describe a project that provided the catalyst for this work. Later sections then introduce ontological research and proceed to use it to lay the foundations for an ontology of information. In the final sections we indicate how such an ontology can be used in the context of emergency management.

Reviewing Statement: This full paper has been fully double-blind peer reviewed for clarity, relevance, significance, validity and originality.

THE PROBLEM OF INFORMATION INTEGRATION

The 2006 Post-Katrina Emergency Management Reform Act requires the US Federal Emergency Management Agency (FEMA) to:

1. Develop a logistics system that provides visibility of logistical assets from procurement to delivery, throughout the entire shipment process.
2. Ensure that existing IT systems are compatible and have the capability to share information.

Although this requirement is specific to FEMA, it should clearly be a goal for logistical asset tracking for any emergency management agency.

In 2008 US Department of Homeland Security (DHS) asked the Office of Inspector General (OIG) to study how well FEMA was managing information technology to support disaster response logistics activities. OIG concluded that such activities were not effectively supported by information systems currently used by FEMA, and that, as a result, FEMA ability to perform timely and effective disaster response might be hindered.

It would be beneficial to bodies involved with information related to emergency management to adopt a common classification scheme for the information types. The most highly developed such scheme is the National Incident Management System (NIMS). NIMS provides a core set of concepts, principles and terminology for incident command and multiagency coordination (US Department of Homeland Security 2008). However, NIMS has not been subject to ontological analysis.

In general, heterogeneity is ubiquitous in emergency management informatics. We have heterogeneous names for entities, process rules, sensor platforms, information systems platforms, data and communication formats, organizations, and even languages. That such heterogeneity can hinder effective disaster response was clearly seen in Haiti following the 2009 earthquake. There is therefore a clear need for an ontology that can provide unified definitions of entities, their properties and relationships, and thus facilitate improved communication in the presence of heterogeneity. The domain of emergency management is complex, and ontological work can provide a framework in which we can specify complex domain entities in terms of primitive components.

SENSOR-BASED INFORMATION SYSTEMS FOR EMERGENCY MANAGEMENT: AN EXAMPLE

Effective emergency management and disaster response can be significantly enhanced by the new sensor-based technologies that are coming on-stream. However, sensor data from many heterogeneous sources, distributed in space and time, can only be effectively utilized if the foundations in informatics and knowledge representation are available to transform the data into information that is timely and meaningful to decision makers.

Emergencies and disasters over the past decade have clearly demonstrated the need for more effective informatic support. In recent history we have seen the failure of information systems to give effective support to emergency management processes. Recent developments in sensor technology provide opportunities for the use of real-time and dynamic information for monitoring and tracking purposes. Given these dramatic changes in the technologies, there is a need to provide in parallel the foundations upon which information systems for emergency management can be developed.

The informatics of emergency management is highly dynamic, with changing scenarios and many mobile entities that need tracking in space and time. In a recent project that was a collaboration between the University of Maine, Global Relief Technologies, and the Maine National Guard, the following two areas of emergency management informatics were considered:

1. **Critical infrastructure monitoring.** Critical infrastructure (e.g., roads, bridges, rivers) is typically stationary. The objective here is to use sensing devices (e.g., remote sensors mounted on satellites, or on-the-ground wireless sensor networks) to monitor important parameters associated with the infrastructure elements. Examples include traffic flows on roads and water levels of rivers and streams. Appropriate monitoring facilitates prediction and response.
2. **Mobile asset tracking.** In response to an emergency, assistance needs to be mobilized. Much of this assistance comes in the form of mobile assets (food supplies, water, generators, emergency shelters). Tracking these assets from supplier to destination facilitates efficient allocation of resources.

The problem area therefore requires the handling of stationary and mobile objects, static and dynamic spatial fields, processes and events, any of which may be defined either over a network or over free space. Networks are particularly important in this area. For example, traffic flow may be modeled as a dynamic field constrained by the road network, or the movement of a truck containing an emergency generator can be modeled as an object whose movement is constrained by a road network. Many critical phenomena such as fires and floods are spread over space, whereas the infrastructure needed for dealing with them is typically restricted to networks (e.g. transport networks, power distribution networks, water distribution networks). Thus in particular we are interested not just in the topology of networks but how they are embedded in space and models of mobility constrained by them. This implies that any ontology of relevant phenomena needs to include relations and interactions between network-bound phenomena and free-space phenomena.

THE ROLE OF ONTOLOGY

In an emergency situation one is confronted not only with information collected by trained personnel acting in an official capacity and reported in accordance with standard procedures and protocols, but also “crowd-sourced” information received as telephone messages, email communications, etc, from members of the general public, as well as increasing amounts of data gathered from automated monitoring devices attached to key infrastructure elements or surveillance vehicles. This information, moreover, may come in many different forms: written reports, oral reports, photographs, sketch maps, numerical measurements, and so on. Bringing together such diverse kinds of information from so many disparate sources presents a major problem to anyone seeking to build up a coherent picture of the developing situation in order to inform rational decision-making, and the development of appropriate informatics support for such activities is thus an important priority.

We believe that ontology has a significant role to play in the general task of improving our ability to handle information. Broadly speaking, there are two distinct reasons for this. Most obviously, ontology can ensure coherent and correct conceptualisation of the real-world domains providing the subject matter of the information to be handled. In emergency management, for example, Di Maio (n.d) refers to a “recurrent lack of consistency”, noting that the terms “displaced person”, “evacuee”, “beneficiary”, and “missing person”, or again, the terms “stock”, “supply”, “intake”, and “donation”, are often used more or less interchangeably. But she also notes that the same is true of terms such as “input”, “data”, and “information”, and this points towards a second role that ontology can play here. This is the second-level task of conceptualising, not the subject matter of the information, but the information itself. While there has been work on using ontologies for handling the domain knowledge representation needs of emergency management (e.g., Segev 2008, Fan and Zlatanova 2010), this additional second-level task has been largely neglected. Even in the development of ontologies for information fusion (e.g., Little and Vizenor 2006) and situation awareness (Matheus *et al.*, 2003), the focus has mainly been on the entities making up the application domain rather than the information entities themselves.

The attempt to construct integrated information systems risks foundering in conceptual confusion unless we can achieve clarity in understanding information itself, providing detailed and coherent answers to such questions as “What is information?”, “What forms does it take?”, “How does it behave?”, and “How is it related to things which are not information?”. Such questions are the province of ontology; in the classic definition of (Gruber, 1993), an ontology is “an explicit specification of a conceptualisation”, the underlying philosophy being that *any* information system designer must at least implicitly conceptualise the domain that is to be handled by the information system, and in order to facilitate the reuse, sharing, harmonisation, and integration of software and data — obviously desirable ends in most cases — it is important to make such conceptualisations explicit.

Gruber’s definition was extended by Studer *et al.* (1998) to include the word ‘formal’ as well as ‘explicit’, implying that the ontology should be machine-readable. This may be particularly important for an information ontology in the emergency management domain. For example, in the Emergency Data Exchange Language EXDL (OASIS, 2006), a *distribution element* includes a set of *content objects*, each of which includes a *content description*; this is a “human-readable text describing the content object” – examples given include “CAP message from FEMA”, “Map of affected area”, and “Photo of missing child”. Use of an information ontology to ground the generation of content descriptions could make them at least partially machine readable, with systematic links to metadata relating to the source, reliability, and other properties of the each content object.

As has been made clear above, it is widely recognised that one of the major obstacles to the creation of adequate informatic support for the emergency management domain is precisely the lack of a sufficiently principled and comprehensive systematisation of the many forms of information, and their complex interrelationships, that such a system must be able to handle. It is in the service of such a systematisation that we propose the development of an ontology of information entities. In the next section we outline the first steps towards the creation of such an ontology. We believe it is important here to begin by stepping back from the immediate concerns of emergency management in order to develop a fully general account of information which will be flexible enough to be adapted to the needs of many different types of emergency management situation.

OUTLINE FOR AN ONTOLOGY OF INFORMATION

An ontology begins with a *taxonomy*, that is, a classification of the entities which form the subject-matter of the domain of interest. Here the term 'entity' is intended to be understood in the broadest possible sense; it may include not only ordinary physical objects, but also (if the ontologist wishes to include them) such things as social institutions, abstractions of various kinds, as well as temporal entities such as processes and events. Each individual entity in the domain is assigned to one or more *classes*, and it is the hierarchical (or other) organisation of the classes that constitutes the taxonomy. The two most fundamental ontological relations are

1. *class subsumption*, generally denoted "is_a", by which one class is said to be a specialisation of another (as, e.g., "horse is_a mammal", indicating that the class of horses is subsumed within the class of mammals — or, as we might say, horses are a special case of mammals); and
2. *class membership*, generally denoted "instance_of", by which an individual entity is said to belong to a given class (e.g., "Bucephalus instance_of horse" which says that Bucephalus is a particular horse).

It is normal for an ontology to begin with a maximally inclusive class, here denoted "Entity", the class of all entities recognised by the ontology. The immediate subclasses of "Entity" may differ from one ontology to another, but a widespread choice is to include a top-level distinction between those entities, such as ordinary physical objects, which endure through time, typically undergoing changes as they do so (known as *continuants* or *endurants*), and those, such as events and processes, which happen or occur (known as *occurrents* or *perdurants*). This two-fold division is exemplified by the widely-used BFO ontology, in which continuants and occurrents form the 'SNAP' and 'SPAN' sub-ontologies respectively (Grenon and Smith, 2004; IFOMIS, 2009).

The distinction between continuants and occurrents is important for our purposes because both types of entity feature in the domain of information entities. To give a simple example, the information that, say, a particular road is flooded at a certain point might be conveyed either by means of a written sentence, which is a continuant, or by means of a spoken sentence, which is an occurrent. Note, however, that the *act* of writing, by which the written sentence comes into being, is an occurrent. Distinct from all of these is the information itself, as opposed to the spoken or written sentences by which it is conveyed.

It is an essential feature of an information entity that it can, in principle, be borne by different physical bearers: information can be *copied*, which means that a new physical bearer is created which carries the same information as some pre-existing bearer. The ontological relation between an information entity and a physical bearer of that entity is one of *dependence*, meaning that the former only exists by virtue of being embodied in the latter. However, this dependence is *generic* because a given information entity is not bound of necessity to a single particular bearer but can be successively (or simultaneously) embodied in many different bearers, which we would normally speak of as different *copies* of the information. Hence Smith's characterisation of an information entity as a *generically-dependent continuant* (Smith, 2009).

To make all this more concrete, consider the case of a member of the public calling the emergency services to report that a bridge has collapsed; the person taking the call makes a written note of the message and passes it to a colleague who enters it into a log held on the computer. Later the log is printed out and photocopies of the printout are distributed to the members of a committee convened to review a developing emergency situation. We see here a series of incarnations of a single information entity — which is, essentially, the proposition that the bridge in question has collapsed — and it is important to note that these incarnations (the bearers of the information) take many forms: a spoken statement, a handwritten note, part of an electronic file, a computer printout, and several photocopies. All of these information bearers may have the same *content*, which is, precisely, the information entity which they all bear.

This information entity is originally generated by an *information origination event*, by which the statement that the bridge has collapsed is first recorded by the member of the public who made the observation. Subsequent propagation of the information in a succession of new bearers is accomplished by a series of *information copying events*. Information origination events and information copying events are classed together as *information events*. (Other types of information event, e.g., information modification events or information merging events, may be defined in terms of these two basic subtypes.) Information events are enacted by *information agents*. In our scenario, these include the member of the public who telephones in as well as the people who receive and write down the message, enter it into the computer, print out the log, and make the photocopies. An information agent is anyone (or anything) that creates a new information bearer, whether this be a first-time bearer for a new piece of information or a new bearer for a piece of information which already exists in another bearer. In many cases an information agent will make use of an *information instrument* to create the new bearer; information instruments include pens, keyboards, printers, and photocopiers.

The picture we have developed so far can be summarised as follows:

- An information entity is borne by one or more information bearers, and exists only so long as it has at least one such bearer.
- The content of an information bearer is the information entity which it carries.
- An information bearer is created in an information event, which may be an information origination event (if the content of the bearer is a new information entity) or an information copying event (if the content of the bearer is derived from some pre-existing bearer).
- An information event is enacted by an information agent, who/which may or may not make use of an information instrument to create the new bearer. The information agent is the source of the information bearer created by the event.

These classes and their relationships are illustrated in Figure 1, in which (as in subsequent diagrams) rectangular boxes are used for classes of continuants, ellipses for classes of occurrents, and rounded rectangles for mixed classes. Continuant information bearers include, for example, written records, computer files, and printouts, while occurrent information bearers include such things as spoken reports and ‘rolling text’ displays.

Broadly similar ideas are presented by Neuhaus and Andersen, who develop an ontology of information for the purpose of describing security aspects of information systems. In (Andersen and Neuhaus, 2009) the emphasis is on the information bearers (‘tokens’), whereas in (Neuhaus and Andersen, 2009), it is rather on the information events, which they call ‘speech acts’. In both these papers, the only kind of content under consideration is *propositional*; this is not the only kind of content we might wish to consider, as will be explained below. The key notions of source and copy which we use here conform closely to the way they are handled in their system.

Information bearers may be classified with respect to both form and content. On the form side, the topmost division is, of course, into continuant and occurrent information bearers. Subclasses of the continuant information bearers include visual (e.g., text, diagrams, maps, pictures, signs), tactile (e.g., Braille), electronic (computer files), as well as various harder to classify types such as gramophone records and CDs. Occurrent information bearers include both auditory (e.g., speech and music, whether live performances or playings of records or CDs) and visual (e.g., rolling text, animations).

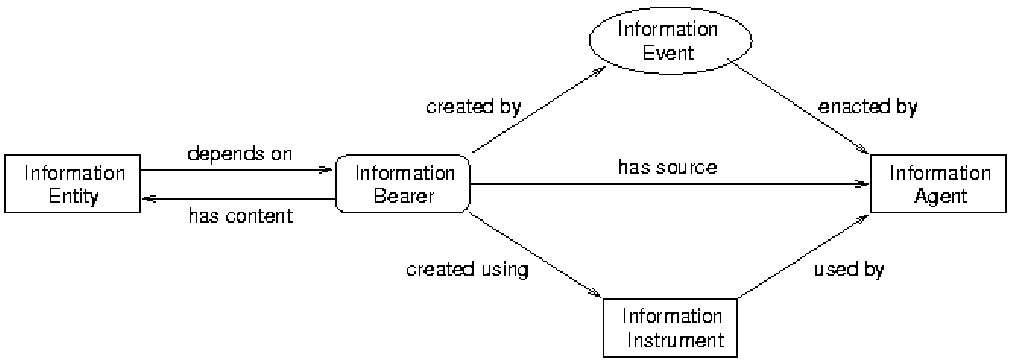

Figure 1. Basic entities and relationships in the information ontology

It is important to note is that content comes in different *levels*. An information bearer may carry content at one level by virtue of content it carries at some lower level. To illustrate, consider a printed text. At a high level, the content of this may be a series of propositions: its *propositional content*. A closely similar content could be borne by a faithful translation of the text into another language. The two texts bear essentially the same propositional content by virtue of bearing distinct *verbal contents*, each consisting of a particular sequence of words. If the text is read out loud, the reading constitutes an occurrent information bearer with the same verbal (and hence propositional) content, but whereas the printed text bears that verbal content by virtue of bearing a particular sequence of printed characters — its *graphemic content*—the spoken text owes it to a sequence of speech sounds, its *phonemic content*. The printed text bears its graphemic content by virtue of consisting of a pattern of visible marks on the paper, this pattern constituting its *visual content*, whereas the computer file from which the printed text was generated has the same graphemic (and hence verbal and propositional) content by virtue of a pattern of electronic activation in the computer’s memory — its *binary content*.

In saying that a bearer possesses some higher-level content by virtue of possessing some lower level content, we do not mean to imply that the lower-level content alone is sufficient to ensure the existence of the higher-level content. Extraction of higher-level content from lower-level content always requires an appropriate context of interpretation. The visual content of a road sign consisting of the symbols “30” inscribed in black on a white background, with a red surround, includes the shapes “3” and “0” juxtaposed, and this content is accessible to any interpreter with appropriate visual faculties. An interpreter familiar with Arabic numerals understands this as a sequence of two numerals – the sign’s graphemic content – from which an interpreter familiar with the decimal system can extract the verbal content, i.e., a representation of the number thirty. Only an interpreter familiar with the conventions of road signage can extract from this the propositional content which, in continental Europe, is something like “The speed limit is 30 km/hr”, whereas in the UK, it is rather “The speed limit is 30 m.p.h.”. We might think of these different contents carried by the same information bearer as different *roles* played by that bearer, depending on the context in which the bearer is considered. This fits in with the view of information objects taken by the D&S extension of the DOLCE ontology (Masolo, Borgo, Gangemi, Guarino and Oltramari, 2003), in which the class of information objects is subsumed within a broader class of *non-agentive functional role*, itself subsumed within the broader class of non-physical endurants.

The hierarchy of levels of content throws light on the traditional distinction between *data* and *information*. Data are often described as ‘raw’, consisting of unprocessed, unstructured, uncontextualised ‘facts’; they are transformed into information by being processed in such a way as to acquire meaning in relation to some wider context. Although this distinction is widely acknowledged and routinely taught in introductory informatics courses, it is not easy to pin down precisely. ‘Rawness’ is a matter of degree: within the hierarchy of content levels, if an information bearer possesses content X by virtue of possessing some lower-level content Y, we might say that Y is the data from which information X is extracted. A single information entity can be simultaneously both data and information, depending on whether it is considered in relation to higher or lower level information entities. In relation to a piece of connected prose in which they occur, words are data; but in relation to the letters of which they are composed, or the printed shapes on the paper, the same words are information. For the information ontology, then, data and information alike consist of information entities.

For many purposes, such as those of emergency managers, the chief value of any information bearer is what it tells us about the world. This may be explicitly encoded in the bearer itself, constituting its propositional content, e.g., a direct report that a certain stretch of road has been rendered impassable by flooding. But information bearers which do not carry explicit propositional content may also be used to support propositions: a picture of the flooded road, which does not in itself carry propositional content, supports the proposition that the road is flooded. Thus another important ontological relation in which information entities may participate is the relation of *support* between an information bearer and the propositions for which it provides confirmatory evidence. Trivially, an information bearer with propositional content supports that content; on top of that, it will also support anything *implied* by the content. A combination of different pieces of information will often jointly support propositions which none of them would support individually: by combining the information that the northern end of a road is impassable with the information that the southern end of the same road is impassible, we might infer that the village situated half-way along the road has become inaccessible. Any record of this inference should include a record of the information entities used to support it; if one of the sources is later discredited, the inference should no longer stand. Neuhaus and Andersen (2009) have begun to develop a formal theory of the informatic support relation; related to this is a substantial body of work on belief revision (Gärdenfors, 1992) which, although not part of the ontology itself, can provide valuable pointers to the way in which the ontology can be used to drive the development of an information system.

Only if there is reason to believe that the information bearer is reliable will the fact that an information bearer supports a particular proposition provide evidence that the proposition is true. This brings us to a consideration of *information quality*, a highly complex topic about which a considerable amount has been written, some from an ontological point of view. Wand and Wang (1996) note that “there does not exist a rigorously defined set of data quality dimensions”, and they attempt to derive such a set from an ontological consideration of the relationship between information systems and the real-world situations they are designed to represent. Frank (2007) regards definition of separable data-quality dimensions as unachievable. Be that as it may, an ontology of information entities is a pre-requisite for an adequate and fully general characterization of different data-quality dimensions. Duckham *et al.* (2001) consider the quality indicators related to the spatial components of data.

Not all measures of information quality apply to all kinds of information. *Correctness*, *accuracy*, and *precision* are generally acknowledged to be important measures of information quality, applying to *representations*, that is, information entities intended to capture properties of some real-world entities (*targets*) they are created to be representations of. Representations include propositions (including measurements, i.e., statements that some measurable quality of a real-world entity has such-and-such value), pictures, diagrams, and maps; they exclude, e.g., pieces of music and non-figurative graphics. A representation may be regarded as picking out a *range* of possible real-world entities: e.g., a picture of a building picks out all possible buildings whose appearance when viewed from some position is compatible with the picture. A representation is correct in relation to a given target if the latter is included in its range; incorrect representations may be more or less accurate depending on how close the range gets to the target – actual measurement of accuracy requires a metric to be established on the relevant domain. The precision of the representation is a measure of how compact the range is: a precise representation picks out a narrower range than a less precise one. A representation may be *unjustifiably* precise, meaning that the process by which the information was generated (e.g., a measurement) was not sufficiently discriminating to pick out such a narrow range; thus another measure of information quality is the extent to which the precision of the representation is justified by the precision of the process. Accuracy and precision are to some extent inversely correlated: the more precise a representation is, the harder it is for it to be accurate; even so, they are clearly distinct measures and thus partially separable. Another dimension, *fidelity*, is a measure of similarity between the contents of two bearers, one of which is derived from the other in an information copying event. Depending on how it is measured, fidelity may also be inversely related to precision, since imprecise representations are intrinsically more tolerant toward copying errors than precise ones; but again these are clearly distinct measures and should not be conflated. In particular, fidelity only applies to the information content of a copy, and then only in relation to the information content of the bearer that it was copied from.

INFORMATION IN THE EMERGENCY MANAGEMENT DOMAIN

With both crowd-sourced data and data collected by sensor networks, we are interested in combining information coming from collections of information bearers. In either case, we need to determine what propositions are jointly supported by a collection of information entities. This is complicated by the fact that the individual information agents are typically spatially distributed, and the information events which they enact are temporally distributed. There are several possibilities, ranging from cases where the spatio-temporal distribution of information events is completely random and unsystematic, with each agent enacting information events independently, to cases where all the elements of a regular array of stationary information agents simultaneously enact information events in a regular sequence. In this section we look specifically at the case of information supplied by sensor networks, but the principles to be expounded should be generalizable to other cases.

A sensor is a kind of information agent, which generates *sensor readings* by enacting *sensor reading events*. A sensor reading is a measurement of a *determinable quality* inhering in some continuant (e.g., temperature of a sample of air or water). Since a sensor determines values only to a certain level of precision, the value it assigns must be interpreted as a range – only with discrete-valued qualities can this reduce to a single value. Thus the determinable quality *has* a value, and the sensor reading *assigns* a range of values to it, the purpose of the sensor being to ensure that the latter includes the former. Our determinable qualities correspond to the *determinands* (or *observables*) of the SensorML modelling language (Open Geospatial Consortium Inc., 2007); these are defined as "parameters or characteristics of a phenomenon subject to observation". An outline for an ontology of sensors and sensor readings, with relations to the general information ontology, is shown in Figure 2.

For gathering information over a wide area it is usual to have an ensemble of separate sensors; these may be in communication with a central monitoring facility, or with each other, forming a *sensor network*. Collectively, they gather a collection of value assignments which may collectively constitute a higher-order information entity relating to some more global state of affairs than anything coming within the purview of a single sensor. Given the important role played by sensor networks and other ensembles of information agents in emergency management, it is important to extend the ontology to handle them. Although the ontological analysis of collectives and ensembles is only in its infancy (Wood and Galton, 2009), some provision is made for handling them in most of the well-known upper ontologies. In BFO (Smith, 2009), for example, there are classes *Object Aggregate* and *Process Aggregate*, covering collections of continuants and occurrents respectively. Both classes (or their equivalents) are needed in the information ontology, the former to include collections of sensors and collections of sensor readings, the latter to include collections of sensor reading events.

A sensor network is a spatial distribution of sensors, which in turn is a collection of sensors. A spatial distribution is more than a mere collection, since the specification of the positions of the sensors forms part of the definition of the former, but not of the latter. Similarly, a sensor network is more than just a spatial distribution, since to define it, one must specify the communication channels that link the sensors together.

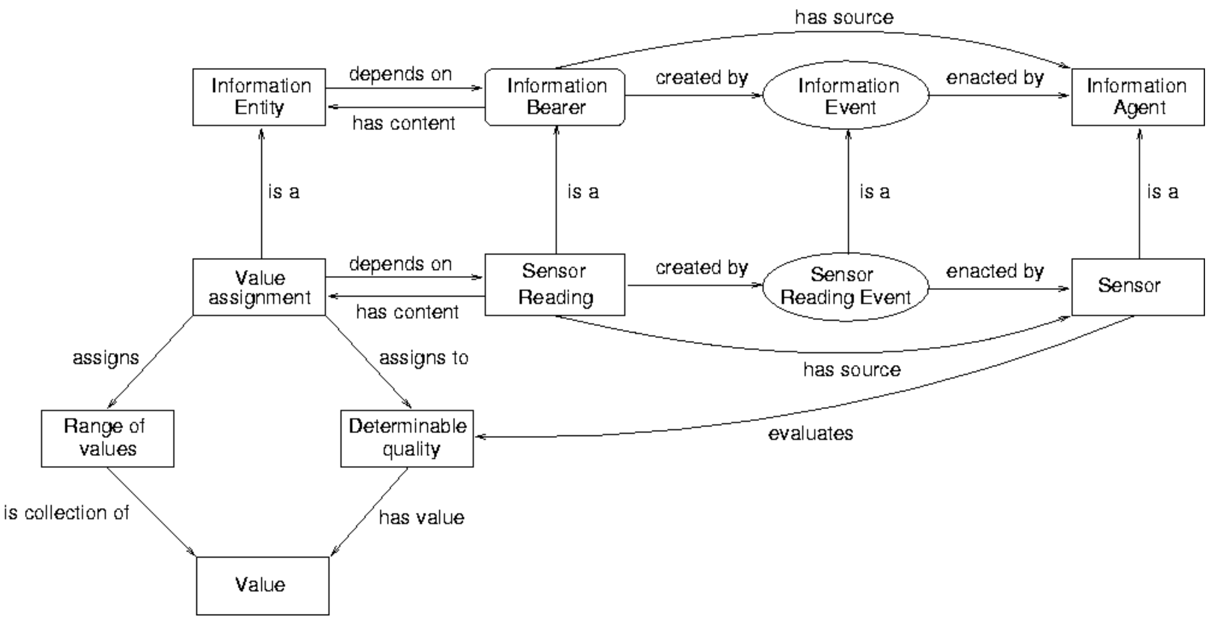

Figure 2. An ontology of sensors and sensor readings as part of the ontology of information entities

A collection of sensors gives rise to a collection of sensor readings. However, there are several different possibilities as regards their spatio-temporal distribution, e.g.,

1. A *time series* of sensor readings, generated successively by a single sensor.
2. A *spatial distribution* of sensor readings, generated simultaneously by an array of sensors.
3. A *spatio-temporal distribution* of sensor readings, generated by an array of sensors over a period.

In Figure 3, the ontology is extended to include collections of sensors; this can be generalised in obvious ways to collections of other kinds of information agents. Note that the class *Time-series of Sensor Reading*s is a continuant, whereas the class *Series of Sensor Reading Events* is an occurrent. The reason for this is that whereas a series of events is something that happens, a complex event built up out of simpler constituent events (in this case the individual reading events), the time-series it creates is a complex information bearer that is built up over a period of time, with each sensor reading event adding a new component to the growing structure.

CONCLUSIONS AND FURTHER WORK

The application of ontological research to the informatics of emergency management (e.g., Di Maio n.d., Segev 2008) has concentrated on conceptualizing the relevant real-world domain knowledge, to the comparative neglect of the information itself. None the less, as we have noted, in the emergency management domain we are confronted with diverse kinds of information relating to diverse kinds of entities, and the task for the emergency manager is to bring all these information entities together in a common framework in order to build up as complete a picture as possible of the situation as it develops (the COP). For reasons of information quality control, it is also necessary to keep track of the provenance of the information entities, and the transformations they undergo as they are transmitted from one bearer to another. Such additional information will be represented as metadata attached to the primary information bearers themselves; key elements of such metadata include

- the immediate source of the bearer,
- the ultimate source of the information it carries,
- some measure of its authenticity or credibility, as determined by (amongst other factors), the reliability of the succession of sources in its history,
- its relation to other information bearers, either corroborative or conflicting, appropriately indexed by the degree of mutual independence between them,
- the propositions it supports, with an indication of the degree of support.

The development of an information system to support such activities must be informed by an appropriate understanding of the different kinds of information that is to be handled, the different forms that the information can take, the transformations it can undergo, and the different levels of accuracy, precision, and trustworthiness that it can manifest. The system must, in short, handle not just information about the real-world entities and events that make up the emergency situation, but also information about that information – meta-information – and it is this that is classified and codified in the information ontology.

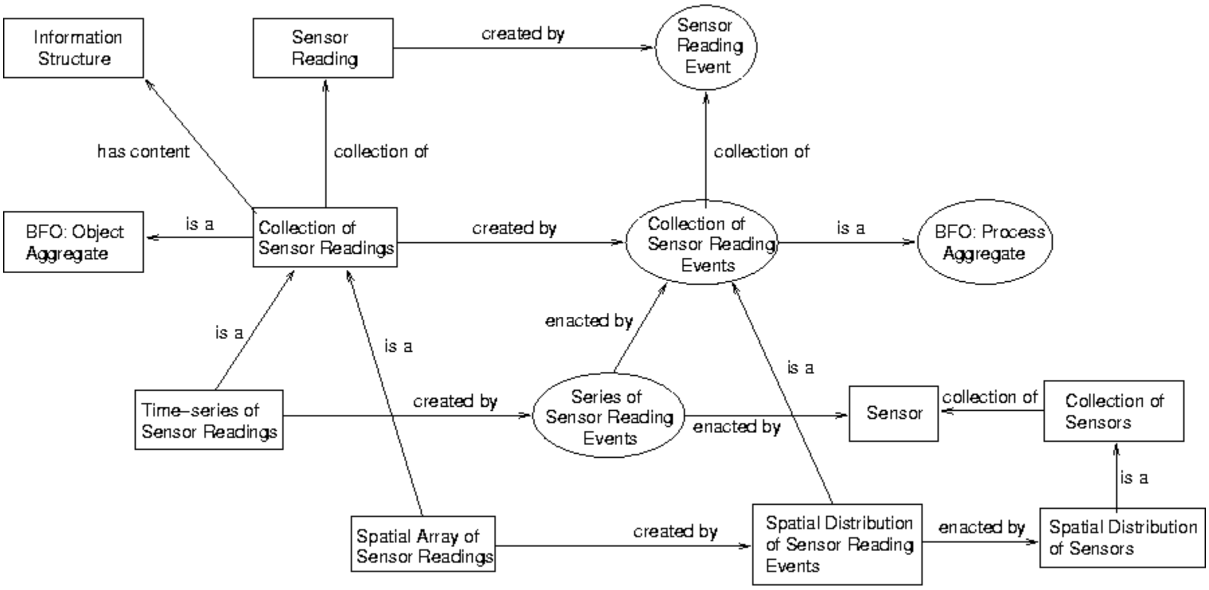

Figure 3: An ontology of collections of sensors and sensor readings

The ontology encodes this knowledge in a purely declarative form; as is generally acknowledged, such separation of the declarative component of a knowledge-based system from its procedural aspects is desirable in order to facilitate knowledge sharing amongst systems which may be incompatible from an operational point of view. One important way in which ontology development differs from object-oriented system design is precisely that the former is entirely free of the procedural considerations which form an essential component of the latter. But the system design may be informed by the ontology in a thoroughgoing way – not least in the determination of which classes the system should comprise and what formal relations should be established between them.

Geospatial information has traditionally been created by government agencies and commercial companies. Recently, however, the widespread availability of GPS, fine-resolution satellite imagery, and mapping software has enabled volunteers to collect and compile accurate geospatial information, and to integrate and disseminate it through the Web, making it freely available. For example, in 2010 the OpenStreetMap project was extremely effective in rapidly producing a detailed street map of Port-au-Prince to support the recovery effort following the Haitian earthquake, and doing so in a matter of days through the actions of "crisis camps" of volunteers. The role of such volunteer driven information will become increasingly important. This is an area where information ontologies can play a role in automating the recruitment and coordination of such volunteers in time-critical situations. A current focus of our work is applying the work reported here on information ontologies to provide the foundation of a system whereby the volunteers provide information that enables them to be linked to specific mapping activities, as well as being connected with each other. Individual information agents are chosen to be part of the information community responding to a specific dynamic, time-critical event based upon factors that include spatio-temporal proximity to the event, domain knowledge, skill set, prior experience with previous information communities, and other potential indicators of trust. The information ontology will also enable integration between human-based and artificial sensor-based information.

Oberle (2006) classified ontologies with respect to the dimensions of purpose, specificity, and expressiveness; in terms of this classification, our information ontology covers a wide range. In purpose, it is both a reference ontology, designed to reduce terminological confusion (cf. the remarks of Di Maio, cited above), and an application ontology used to drive an application; and in relation to specificity, it includes a generic component covering such broad categories as information bearer, information entity, and information event, but also serves as a domain ontology, having reference to specific types of information agent (e.g., people, sensors), specific types of information bearers (e.g., maps, pictures, telephone messages, sensor readings), and so on.

In future work we will extend the information ontology to incorporate all the different kinds of information entities, information bearers, and information events that are of importance to the emergency management domain, and then we will link the information ontology to relevant domain ontologies handling some specific types of emergency situation, e.g., flooding, fires, or terrorist activities. The links between the ontologies will cover the relationships between information entities and the real-world objects or processes that they are about, and also relationships arising from the fact that information agents and even information bearers may already be referenced in the domain ontology (e.g., people affected by an incident can act as information agents).

REFERENCES

1. Andersen, W. and Neuhaus, F. (2009) An ontological approach to information access control and provenance, *Proceedings of OIC 2009 (Ontology for the Intelligence Community)*, George Mason University, Fairfax Virginia Campus, 21st–22nd October 2009. http://c4i.gmu.edu/OIC09/papers/OIC09 7 AndersenNeuhaus.pdf.
2. Duckham, M., Mason, K., Stell, J. and Worboys, M.F. (2001) A formal approach to imperfection in geographic information, *Computers, Environment and Urban Systems*, 25: 89-103.
3. Fan, Z. and Zlatanova S. (2010) Exploring ontology potential in emergency management, *Proceedings of the Gi4DM Conference - Geomatics for Disaster Management,* February 2010, Torino, Italy.
4. Frank, A. U. (2007) Incompleteness, Error, Approximation, and Uncertainty: An Ontological Approach to Data Quality, *Geographical Uncertainty in Environmental Security*, NATO Security through Science series, 107-131.
5. Gärdenfors, P. (1992) *Belief Revision*. Cambridge University Press, Cambridge UK.
6. Grenon, P. and Smith, B. (2004) SNAP and SPAN: Towards dynamic spatial ontology, *Spatial Cognition and Computation*, 4, 1, 69–104.
7. Gruber, T. R. (1993) A translation approach to portable ontology specifications, *Knowledge Acquisition*, 5, 2, 199–220.
8. IFOMIS (2009) Basic formal ontology (BFO), http://www.ifomis.org/bfo, Accessed 15/10/2010.
9. Di Maio, P. (n.d.) An open ontology for open source emergency response system. http://opensource.mit.edu/papers/TOWARDS AN OPEN ONTOLOGY FOR ER.pdf (accessed 15/7/10).
10. Masolo, C., Borgo, S., Gangemi, A., Guarino, N. and Oltramari, A. (2003) *WonderWeb deliverable D18*. Technical report, Laboratory for Applied Ontology, ISTC-CNR, Trento.
11. Matheus, C. J., Kokar, M.M., and Baclawski, K., A core ontology for situation awareness, *Proceedings of the Sixth International Conference on Information Fusion*, 2003, 545-552.
12. Neuhaus, F. and Andersen, W. (2009) The bigger picture: Speech acts in interaction with ontology-based information systems. In Okada, M. and Smith, B., editors, *Interdisciplinary Ontology: Proceedings of the Second Interdisciplinary Ontology Meeting, February 28th–March 1st, 2009, Tokyo*. Keio University, Open Research Centre for Logic and Formal Ontology, 45-56.
13. OASIS (2006) Emergency Data Exchange Language (EDXL) Distribution Element v. 1.0, OASIS Standard EDXL-DE v1.0, May 2006. http://docs.oasis-open.org/emergency/edxl-de/v1.0/EDXL-DE_Spec_v1.0.pdf
14. Open Geospatial Consortium Inc. (2007) OpenGIS Sensor Model Language (SensorML) Implementation Specification, Technical Report.
15. Oberle, D. (2006) *Semantic Management of Middleware*, Volume I of *The Semantic Web and Beyond*, Springer, New York.
16. Segev, A. (2008) Adaptive ontology use for crisis knowledge representation. In *Proc. of the 5th Int. ISCRAM Conf.*, ed. Fiedrich, F. and van de Walle, B. Washington, DC., 285-293.
17. Smith, B. (2009) Basic formal ontology, http://icbo.buffalo.edu/Presentations/Smith.ppt.
18. Studer, R., Benjamins, R. And Fensel, D (1998) Knowledge engineering: Principles and methods. *Data and Knowledge Engineering*, 25(1-2), 161-198.
19. US Department of Homeland Security (2008) *National Incident Management System*. http://www.fema.gov/pdf/emergency/nims/NIMS core.pdf, December 2008.
20. Wand, Y. and Wang, R. Y. (1996) Anchoring Data Quality Dimensions in Ontological Foundations, *Communications of the ACM*, 39, 11, 86-95.
21. Wood, Z. and Galton A. (2009) A Taxonomy of Collective Phenomena, *Applied Ontology*, 4, 3-4, 262-292.