Defence Research and Development Canada	Recherche et développement pour la défense Canada

DEFENCE R&D DÉFENSE

Theoretical foundations and proposed applications of Threat Ontology to information fusion

E. Little
G. Rogova
Calspan-UB Research Center, Buffalo

A.-C. Boury-Brisset
DRDC Valcartier

Defence R&D Canada – Valcartier
Technical Report
DRDC Valcartier TR 2005-269
November 2008

Canada

Theoretical foundations and proposed applications of Threat Ontology to information fusion

E. Little

G. Rogova
Calspan-UB Research Center, Buffalo

A.-C. Boury-Brisset
DRDC Valcartier

Defence R&D Canada - Valcartier

Technical Report

DRDC Valcartier TR 2005-269

November 2008

Author

Anne-Claire Boury-Brisset

Approved by

Éloi Bossé

Section Head – Decision Support Systems

Approved for release by

Christian Carrier

Chief Scientist

Abstract

Information Fusion is the process of utilizing information from various sources to produce knowledge about objects and situations in order to assist humans in decision making. It is argued here that fusion processes would be better equipped to produce these entities estimates if designed in conjunction with formal ontologies. Such ontologies can provide fusion system designers with a comprehensive, and computationally tractable, phenomenological description of a given domain. A properly conceived formal ontology is necessary to provide a structure for analyses of domain-specific objects, object attributes, and relations and to assure both interoperability and reusability of the designed fusion system for different domains. In this sense, a formal ontology *informs* the fusion process by providing a phenomenological description of a given state of affairs. At the same time the fusion process *constrains* the formal ontology in terms of its size, scope, and relevance for certain decision-making needs. This report presents ontology related concepts for building a threat ontology (ThrO) to be used in various information fusion applications involving threats.

Résumé

La fusion d'informations est le processus d'intégration de données provenant de sources multiples en vue de produire des connaissances relatives à des objets et des situations afin d'aider les humains à prendre des décisions éclairées. Dans ce rapport, nous soutenons l'idée que les processus de fusion seraient mieux équipés pour produire ces estimations s'ils étaient conçus de concert avec des ontologies formelles. De telles ontologies peuvent fournir aux concepteurs de systèmes de fusion une description phénoménologique, globale, exploitable automatiquement, d'un domaine donné. Une ontologie formelle conçue rigoureusement est nécessaire afin de fournir une structure d'analyse des objets du domaine, leurs attributs et relations, et assurer l'intéropérabilité et la réutilisation de systèmes de fusion dans différents domaines. En ce sens, une ontologie formelle informe le processus de fusion en fournissant une description phénoménologique d'un état donné d'affaires. En même temps, le processus de fusion contraint l'ontologie formelle en matière de taille, portée et pertinence pour certains besoins de la prise de décision. Ce rapport présente les concepts et les étapes méthodologiques pour la construction d'une ontologie formelle relative au concept de menace pouvant être utilisé dans le cadre d'applications de fusion d'informations variées.

This page intentionally left blank.

Executive summary

Information Fusion is the process of utilizing information from various sources to produce knowledge about objects and situations in order to assist humans in decision making. Recent research efforts in information fusion focus on the design of fusion systems that address high levels of information fusion, i.e. situation and threat assessment. Scientists at DRDC Valcartier have initiated research in order to contribute in advances in the domain and make progress in the building of situation and threat assessment support systems. In this context, a research collaboration between DRDC Valcartier and the Center for Multisource Information Fusion (CMIF) at the University of New York in Buffalo has investigated the theoretical foundations of threat ontology for information fusion with the intent of building ontologically grounded fusion systems.

This report presents results from this collaboration. It is argued here that fusion processes would be better equipped to produce these entities estimates if designed in conjunction with formal ontologies. Such ontologies can provide fusion system designers with a comprehensive, and computationally tractable, phenomenological description of a given domain. A properly conceived formal ontology is necessary to provide a structure for analyses of domain-specific objects, object attributes, and relations and to assure both interoperability and reusability of the designed fusion system for different domains. In this sense, a formal ontology *informs* the fusion process by providing a sufficiently complex phenomenological description of a given state of affairs. At the same time the fusion process *constrains* the formal ontology in terms of its size, scope, and relevance for certain decision-making needs.

The document describes ontological related concepts for the building of a Threat Ontology that may be used in various information fusion applications. First, it presents the information fusion domain and the way formal ontologies and information fusion may interact. Then, chapter 3 introduces the philosophical underpinnings of basic formal ontology, in particular with the representation of upper-level, abstract categories of space and time. In chapter 4, it the concept of a threat ontology with its components and properties are described. The methodological approach to ontology building is illustrated with a case study of a dirty bomb. Finally, some conclusions are provided.

Little, E., Rogova, G., Boury-Brisset, A.-C. 2008. Theoretical foundations and proposed applications of threat ontology to information fusion. DRDC Valcartier TR 2005-269. Defence R&D Canada.

Sommaire

La fusion d'informations est le processus d'intégration de données provenant de multiples sources en vue de produire des connaissances relatives à des objets et des situations afin d'aider les humains à prendre des décisions éclairées. Les efforts de recherche récents en fusion d'informations visent à concevoir des systèmes de fusion qui traitent des niveaux de fusion supérieurs, à savoir l'évaluation de la situation et de la menace. Plusieurs scientifiques à RDDC Valcartier ont entrepris des travaux de recherche afin de contribuer aux avancées dans ce domaine et à construire des systèmes de fusion supportant ces niveaux. Dans ce contexte, une collaboration de recherche entre RDDC Valcartier et le centre de fusion multi-sources de l'Université de New York à Buffalo (CMIF) a investigué les fondations théoriques du concept d'ontologie de la menace dans le cadre de la fusion d'informations dans le but de bâtir des systèmes de fusion fondés sur des bases ontologiques.

Ce rapport présente les résultats de cette collaboration. Nous soutenons l'idée que les processus de fusion seraient mieux équipés pour produire ces estimations s'ils étaient conçus de concert avec des ontologies formelles. De telles ontologies peuvent fournir aux concepteurs de systèmes de fusion une description phénoménologique, globale, exploitable automatiquement, d'un domaine donné. Une ontologie formelle conçue rigoureusement est nécessaire afin de fournir une structure d'analyse des objets du domaine, leurs attributs et relations, et assurer l'intéropérabilité et la réutilisation de systèmes de fusion dans différents domaines. En ce sens, une ontologie formelle informe le processus de fusion en fournissant une description phénoménologique suffisamment formelle d'un état donné d'affaires. En même temps, le processus de fusion contraint l'ontologie formelle en matière de taille, portée et pertinence pour les besoins de la prise de décision.

Ce rapport présente les concepts et les étapes méthodologiques pour la construction d'une ontologie formelle relative au concept de menace pouvant être utilisé dans le cadre d'applications de fusion d'informations variées. Nous introduisons les bases philosophiques soutenant le concept d'ontologie formelle, en particulier par la représentation de catégories abstraites de haut niveau relatives aux notions de temps et d'espace. Puis, nous décrivons l'application de ces concepts à une ontologie de la menace. L'approche méthodologique est illustrée par un exemple de « bombe sale ». Finalement, des conclusions sur ce travail de recherche sont formulées.

Little, E., Rogova, G., Boury-Brisset, A.-C. 2008. Theoretical foundations and proposed applications of threat ontology to information fusion. DRDC Valcartier TR 2005-269.
R et D pour la Défense Canada.

Table of contents

List of figures

This page intentionally left blank.

1. Introduction

This report represents an initial attempt to provide a theoretically grounded methodology for defining threats of various types, which, in turn, could be used in information fusion applications. A comprehensive and accurate understanding of threat can be achieved through the construction of a formal Threat Ontology (ThrO). Formal ontologies are becoming increasingly widespread in the computer science and artificial intelligence (AI) communities. The ontology discussed here bears certain similarities to other types of computational ontologies found in the literature (Gruber, 1995; Lenat and Guha, 1990; Kokar and Wang, 2002; Menzel and Mayer, 1990; Uschold and Gruninger, 1996; Welty and Smith, 2001; Guarino, 1998; Noy & McGuinness, 2001). However, it differs in the sense that it begins with a more thorough treatment of the metaphysical structure of threat ontology, represented by upper-level, abstract categories of space and time in general (*SNAP* and *SPAN*), considered initially as somewhat independent from any domain-specific applications. In this sense, the Threat Ontology was constructed from both a top-down philosophical perspective (abstract level → domain-specific level) and a bottom-up application-based engineering perspective (domain-specific level → abstract level). Conversely, many computationally driven ontologies stemming from the computer science and AI communities begin their investigations from the lower levels of domain-specific items and define upper levels as needed on an ad hoc basis.

The Threat Ontology represents a somewhat different approach to ontology construction; one which attempts to provide both a deeper philosophical understanding of threats, their features and attributes, their parts, their relations, as well as a broader, application-based understanding of threats pertaining to the various needs of fusion system designers and users. Since providing both depth and breadth of knowledge is the goal of ontology construction, it is important therefore to have both a philosophically sound as well as domain-relevant ontology product.

The document describes ontological related concepts for the building of a Threat Ontology that may be used in various information fusion applications. First, it presents the information fusion domain and the way formal ontologies and information fusion may interact. Then, chapter 3 introduces the philosophical underpinnings of basic formal ontology, in particular with the representation of upper-level, abstract categories of space and time. In chapter 4, it the concept of a threat ontology with its components and properties are described. The methodological approach to ontology building is illustrated with a case study of a dirty bomb. Finally, some conclusions are provided.

2. Information fusion and ontology

In the military context, information fusion is defined as “a process dealing with the association, correlation, and combination of data and information from single and multiple sources to achieve refined position and identity estimates, and complete and timely assessments of situations and threats, and their significance” (White, 1988). Information fusion originated in defence research but recently is becoming more established in many non-military applications such as signal processing, robotics, transportation, remote sensing, optical character recognition, medical decision making, and crisis management (Llinas and Waltz, 1990; Hall and Llinas, 2001; Shimshoni and Intrator, 1998; Abidi. and Gonzalez, 92; Rogova and Llinas, 1996; Benediktsson and Kanellopoulos, 1999; Rogova, 1994; Rogova and Stomper, 2002; Scott and Rogova, 2004).

There are various fusion processing models considered in the literature (White, 1988; Dasarathy, 1997; Bedworth and O'Brien, 2000; Salerno, 2002; Endsley 1995; Blasch and Plano, 2003). To help characterize the primary functions involved in information fusion, the Joint Directors of Laboratories (Steinberg et al. 1998) proposed a functional model known as the JDL model. The process at the lower level of abstraction, corresponding to the JDL Levels 0/1, operates with numerical data (measurements, features) and employs numerical, algorithm-oriented methods free of contextual significance. This process produces information about location, kinematics, and identity of single objects. Information here can be defined as judgments made based on data for resolving questions related to data uncertainty model, disclose or reveal distinctions between items, or enable new actions (Ören, 2001). At a higher level of abstraction, corresponding to the JDL Levels 2/3 (situation and threat assessment), this information, along with historic databases and expert knowledge, is used to provide decision makers with a contextual understanding and interpretation of current and future events and behaviors of interest. Processes at this higher level of abstraction operate with symbols or belief values for context processing and employ both numeric and symbolic techniques such as real-time knowledge-based systems, evidence theory, logic, belief networks, and neural networks among others. The results of fusion, at any level of abstraction, are continually evaluated to define the needs for additional sources, or modification of the process itself, to achieve better results (Level 4 of the JDL model). The product of information fusion is a stored dynamic representation of relations between objects and events obtained through fusion (Lambert, 2003), enabling effective action in a corresponding domain. The information fusion process described above is shown in Figure 1.

A considerable amount of focus over the past 10 years has been in research and applications related to numerical processing of data for object assessment. As a result, steady progress is being made in this area. Fusion technology research and development associated with contextual understanding and interpretation of current and future events and behaviors has not been given as much attention. Thus its capabilities are still evolving (Tangney, 2002) and new approaches to situation assessment and prediction are required.

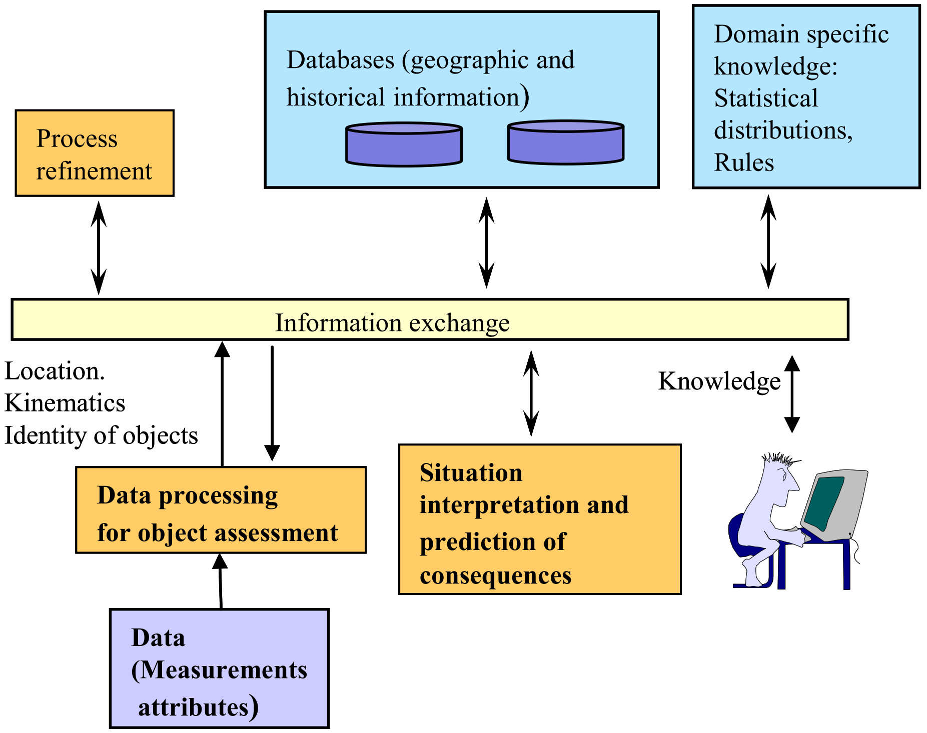
Process refinement
Databases (geographic and historical information)
Domain specific knowledge: Statistical distributions, Rules
Information exchange
Location. Kinematics Identity of objects
Knowledge
Data processing for object assessment
Situation interpretation and prediction of consequences
Data (Measurements attributes)

Figure 1. *Information fusion process based on the JDL model*

Generally, the processes of situation assessment and consequence prediction can be described as follows. Let possible relevant situations or states of the real world (states of situations) ω_j, j = 1,…,J, be elements of a finite set Ω: $\omega_j \in \Omega$. The current state of a decision maker's knowledge about relevant states of the real world (the agent's epistemic state) is described by a subset $K \subset \Omega$ containing an actual or predicted situation. We can also define a belief structure over a set of Ω (represented by a likelihood function $d\in [0,1]$: $d: \Omega \rightarrow [0,1]$. The form of the likelihood function is defined by the uncertainty theory considered, e.g., represented by probability within the Bayesian framework (Ash, 1970), possibility or necessity in the framework of possibility theory (Dubois and Prade, 1988), or by plausibility or belief within evidence theories (Shafer, 1976; Smets and Kennes, 1994). A body of knowledge of a decision-making agent can be viewed as $K = \{U \subset \Omega \mid d(U) \neq 0\}$ (a set of possible current or predicted situations). As in (Halpern, 2003) we can call a tuple (Ω,K,d) an *epistemic belief structure*. In a dynamic world, the state of knowledge of a decision-making agent and a likelihood function over possible situations depend on time: $K = K(t)$ and $d = d(t)$. The likelihood $d(t)$ is defined as a function of attributes of possible worlds, relations between these attributes, and a priori knowledge about these attributes and relations.

The objective of information fusion at higher levels of abstraction is to utilize incoming information at time t along with knowledge about a situation at time t – Δt, historical databases, and expert knowledge to produce a new state of the epistemic belief structure (Ω, K(t), d(t)), for characterizing situations at time t or (Ω, K(t + Δt), d(t + Δt)) for prediction of situations at time t + Δt.

The goal of information fusion is to exploit all available information and prior knowledge to improve an agent's knowledge about the real world by building an adequate epistemic belief structure. In order to be able to reason about and produce an adequate representation of the states of the real world, the fusion process needs a formally structured and computationally tractable representation of related types of objects, events, and relations between them. So it would be advantageous for fusion systems to be ontologically based, since then they would have a guaranteed correspondence with features of the real, objectively independent world.

The way in which formal ontologies and fusion systems interact is via a three-place relation of informing and constraining (see Figure 2). The ontology provides fusion with a coherent and consistent phenomenological description of reality, but at the same time this information is constrained by the needs and requirements of the fusion processes and its relevance to some decision maker's (i.e., user's) tasks or objectives. Constraint on the ontology is necessary to limit the scope of ontological investigation to a specific domain of interest. Also, since formal ontologies are primarily designed by hand, one needs to consider time, effort and cost of their construction. Lastly, constraining an ontology based on needs and requirements of the system can allow one to avoid providing an overabundance of superfluous information within the ontology, which can lead to processing lags in real-time computations.

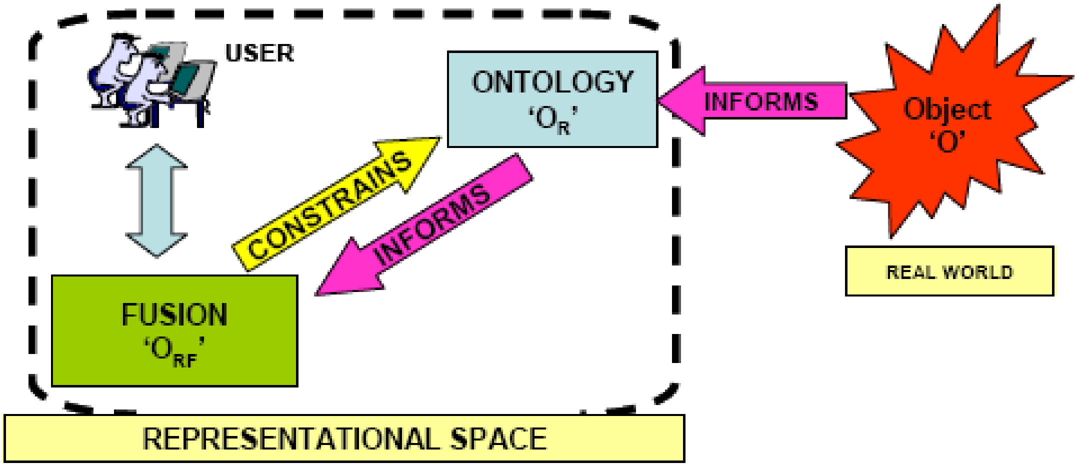

Figure 2. *Informing-constraining relation between ontology and fusion*

Some object (O) (or event (E) in certain circumstances) existing in the world serves as information for the ontology (both upper- and lower-levels). The ontologist constructs a formal *representation* (O_R) of O, which accurately describes certain properties, attributes, and behaviors of O, as well as relations between O and other items. The ontology then informs the fusion node of those formally structured properties, attributes, behaviors and relations of O_R to be processed algorithmically by the fusion system, producing a *fused representation* of O (O_{RF}). Simultaneously, the fusion node is constraining the information within the ontology based on its needs and requirements for a certain goal (or set of goals). The fusion system outputs its fused state estimate (O_{RF}) to either a fully or semi-automated system or a human decision-maker. Because this entire process is ultimately informed by reality (as all properly transparent, realist ontologies are), the relations between the epistemically driven processes of users and the ontologically based structure of reality can also be appropriately understood and modeled if needed. In all, this lends an ontology of the real, independently existing world, which is specifically suited for a user's comprehension of a given portion of it.

One of the main challenges of designing fusion systems is related to the problem of the consistent and comprehensive representation of domain specific types of objects and situations, and the relations between these entities and the environment (Roy, 2001; Roy et al. 2002). Formal ontologies aim at providing fusion with a sufficiently comprehensive model of reality, containing both a well-defined upper-level categorical structure (the most basic items and relations) as well as a sufficiently comprehensive lower-level categorical structure (the individual, domain-specific items and relations). A formal ontology of this sort would present a taxonomy of qualitative features of the world. We hypothesize that this type of qualitatively based taxonomy, along with established thresholds/boundaries of selected features of the world, could produce a computationally tractable framework for improved fusion system design.

Moreover, a formal ontology containing a metaphysically grounded upper-level category structure, not designed for a domain-specific purpose, can provide a general structure that is applicable to numerous domains, since upper-level items are highly abstract and can therefore subsume any and all specific items existing in a given empirical domain. A fusion system that is ontologically structured could be reusable across various specific domains. In fact, the ontology may even help in eliciting new forms of knowledge discovery, since there may be patterns of similarity found between disparate wings of domain-specific ontologies when considered in terms of their upper-level (abstract) ontological structure. In addition, all fusion processes designed in accordance with a upper-level formal ontology are able to operate in concert, which simplifies certain problems of system integration.

The quality of knowledge produced by an ontology-based fusion system is defined by the quality of the ontology used to build this fusion system, which in turn is defined by the level of accuracy, with which the ontology represents reality. Therefore, a very important question here is how to validate the adequacy of an ontology to the domain it represents in order to understand whether it provides a required consistent and comprehensive representation of the environment.

3. Fundamental ontology

The study of metaphysics can be divided into two distinct branches: the first being ontology and the second being epistemology.

Ontology is the study of what is, what exists, what can be logically categorized. Ontologists attempt to capture the most basic structures of reality by developing accurate and comprehensive formal systems that transparently model existing places, times, entities, properties, and relations. The transparency of an ontology is important, because the goal of ontology construction should be to model items as they stand in the world (Smith 1989; Smith and Grenon, 2003; Little, 2002, 2003). In this sense, it should not be limited to issues such as how items are perceived, by whom they are perceived, or how they are conceptually understood/reported upon. An ontology that is non-transparent may give an inaccurate portrayal of reality, in turn, presenting a skewed, perceptively based representation of it. In such cases, the ontology in question could provide a fallacious model of reality, one that is not supported by empirical facts or even common sense. In this sense, ontologists should operate like an empirical scientist, meaning that they should attempt to provide an accurate, third-person-based, independently observable description of the world, distinct from cultural, linguistic, or other types of cognitive biases.

The second branch of metaphysics, epistemology, deals with theories of knowledge and the mental operations of agents (i.e., knowers) who are their bearers. Epistemology is unlike ontology, in that it is unable to provide a third person-based, independently observable description of the world (see Figure 3). Knowledge is always tied to some agent who possesses it. Thus, the history of philosophical studies is strewn with thinkers such as Descartes, Berkeley and Kant, all of whom argued, to various degrees, that reality is necessarily observer-dependent. According to theories such as theirs, the existence of reality is always suspected, since it can never be objectively verified, only subjectively verified. Therefore, reality has been understood by many traditional epistemologists as a residual byproduct of conscious states.

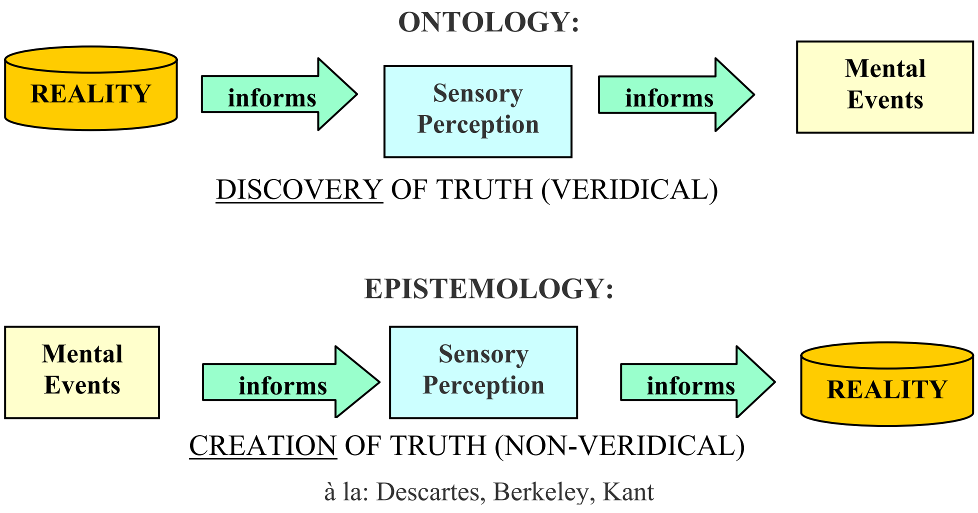

Figure 3. *Distinction between ontology and epistemology*

Because minds are products of the neurological functions of brains, and the latter are part of the objective, ontological furniture of the world, one can argue that all epistemic items are necessarily dependent on some physical substrata or other (Searle, 1992; Dennett, 1991; Little, 2002). Though the framework of an ontology is the product of human creation, and thus a product of certain epistemic processes involving formal logic, semantics, etc., its transparency guarantees its accuracy and correspondence to an external, mind-independent reality.

3.1 Basic Formal Ontology (BFO)

The Basic Formal Ontology (BFO) is a theory, which is currently under development at the Institute for Formal Ontology and Medical Information Sciences (IFOMIS) within the University of Leipzig and at the University of Buffalo (Grenon 2003a, Grenon 2003b). The BFO project is aimed at providing a fundamental theory of ontology, which is not only realist, but also perspectivalist, fallibalist, and adequatist (Grenon & Smith, 2003).

Perspectivalism is the view that there may be various legitimate perspectives on the world, drawn from different individuals, but which are all constrained by the general theory of realism. This means that different people may, and often do, have different perspectives on the world, but these perspectives must be supported by real evidence, as is the case in empirical experiments. *Fallibalism* is the position that theories, systems, models, etc., are works in progress, meaning that they are always open to further scrutiny and can be subject to revision. *Adequatism* is the view that philosophical reductionism is incorrect. There is not necessarily one overarching item, or set of items, to which everything within a domain can be reduced.

In order to maintain its realist, perspectivalist, fallibalist, and adequatist position, BFO should be understood as a general metaphysical ontology from which various regional ontologies can be constructed. BFO is an upper-level theory of ontology. It is a rationally driven, philosophical exercise, which attempts to describe a system's metaphysical structure somewhat independent of its specific application or use-value. BFO serves as an overarching template through which numerous orthogonal, regionally specific ontologies can be created. In this sense, BFO provides the philosophical underpinnings for specified engineering models. It assures that various models will be consistent, reusable and interoperable, since it provides a common, upper-level metaphysical structure, which is consistent with the overall structure of reality (Little, 2003).

It has been effectively argued by Grenon and Smith (2003) that a proper formal ontology cannot confuse continuants (objects) with occurrents (processes, events, changes, activities). These terms can be defined as follows:

Continuant (df.) = entities that have continuous existence and endure through time in spite of various rates of change (e.g., a human organism, the paint on your car, a tree).

Occurrent (df.) = Four dimensional items which occur purely within time and unfold themselves through some period of time. (e.g., a musical measure, the transmission of serotonin between two neurons, waving to a friend).

If one thinks for a moment about the heart, one realizes that at any given moment there is both a spatial entity, which is responsible for pumping blood throughout the body, and at the same time, there is a function of that item which is the process, or string of processes, of the pumping activity. One would be mistaken to confuse the object 'heart' with its processes of 'pumping,' 'squeezing,' 'relaxing,' etc. Objects and processes possess different properties, different part-relations, and different frames of reference. I can, for instance, remove someone's heart from one's body, place it on a table, dissect it, and put it back. It is by all accounts a part of someone's body, which can be removed, replaced, or exchanged. But what about the heart's function, that is, its occurrence of pumping blood? I cannot lay that on the table, nor can I exchange it, or physically dissect it. The function of the heart is not part of the heart (or the body) the way in which a valve is part of the heart. The function of the heart's pumping is a temporal process, which the heart, as a spatial object, undergoes through time. A person's heart, like so many other objects, can maintain its identity throughout various temporal changes.

The distinction between continuants and occurrents exists in order to avoid certain traditional philosophical problems of identity, where it has been argued that if something changes any of its parts, it becomes a new object (i.e., it obtains a new identity). If this is the case, then every time someone takes a shower, thereby sloughing off skin cells, they become a new person; every time someone drives their car, thereby losing molecules of rubber from the tires and paint from the body, they possess a new car. If one is serious about problems of identity, then normally one wishes to say that a given object can maintain its identity, from moment to moment, even though it undergoes subtle changes.

3.2 SNAP and SPAN

Given the importance of distinguishing between occurrents and continuants, there is a need for a unified ontological framework, which can treat both independently, while at the same time, speak to their interconnection. Therefore, the BFO is composed of two orthogonal sub-ontologies, SNAP and SPAN, which share transcategorial relations to one another. SNAP is used to categorize continuants (spatial entities), whereas SPAN is used to categorize occurrents (temporal items). The SNAP portion of an ontology can be thought of as a snapshot of reality, where there is no temporal progression. It represents a strictly three-dimensional ontology in which only continuants (i.e., objects, properties or attributes of objects, and a-temporal relations between them) are represented. The basic structure of SNAP can be seen in Figure 4.

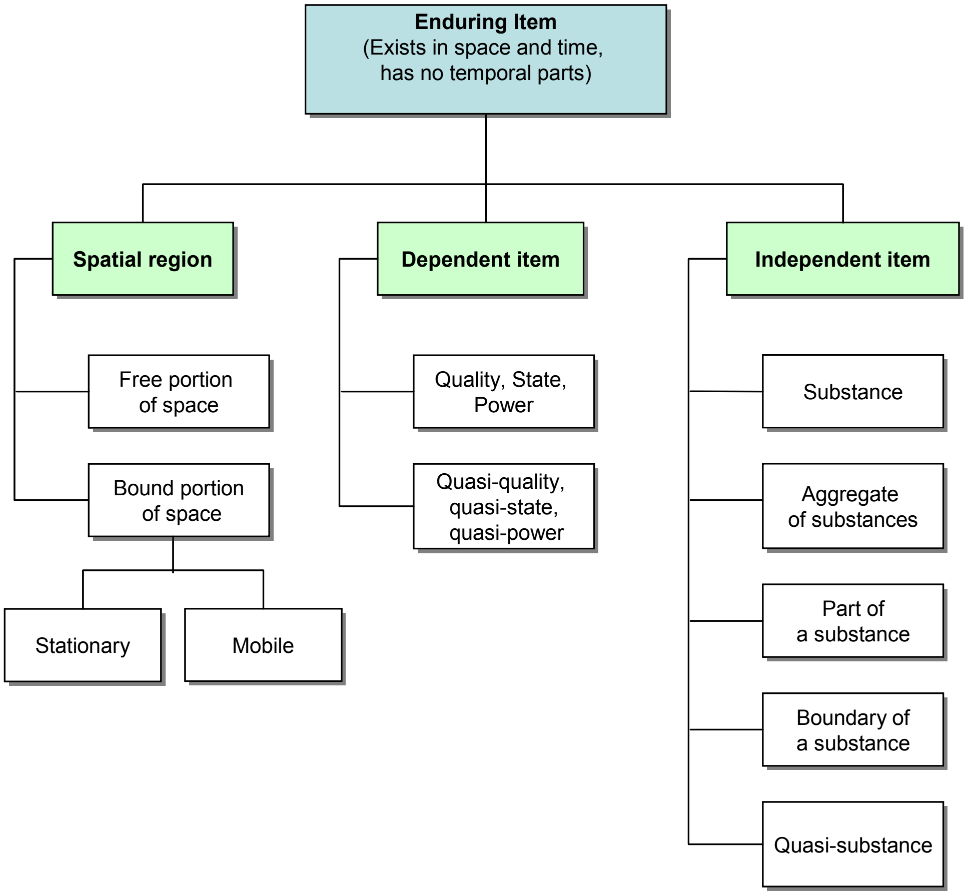

Figure 4. *SNAP BFO upper-level categories*

The SPAN portion of an ontology on the other hand can be thought of as a videoscopic representation of reality, where only occurrents (i.e., events, processes, activities) are represented. SPAN ontologies are used to represent the unfolding of processual entities (see Figure 5). The entities, which are contained within a SPAN ontology, are not as crisply defined as entities within a SNAP ontology, since any processual entity possesses an intrinsic complexity in the sense that it exists as a duration, temporal extension, or succession. This means that it always contains a beginning, a middle and an end, though the boundaries between these temporal phases are normally fiat (i.e., are the process of human deliberation).

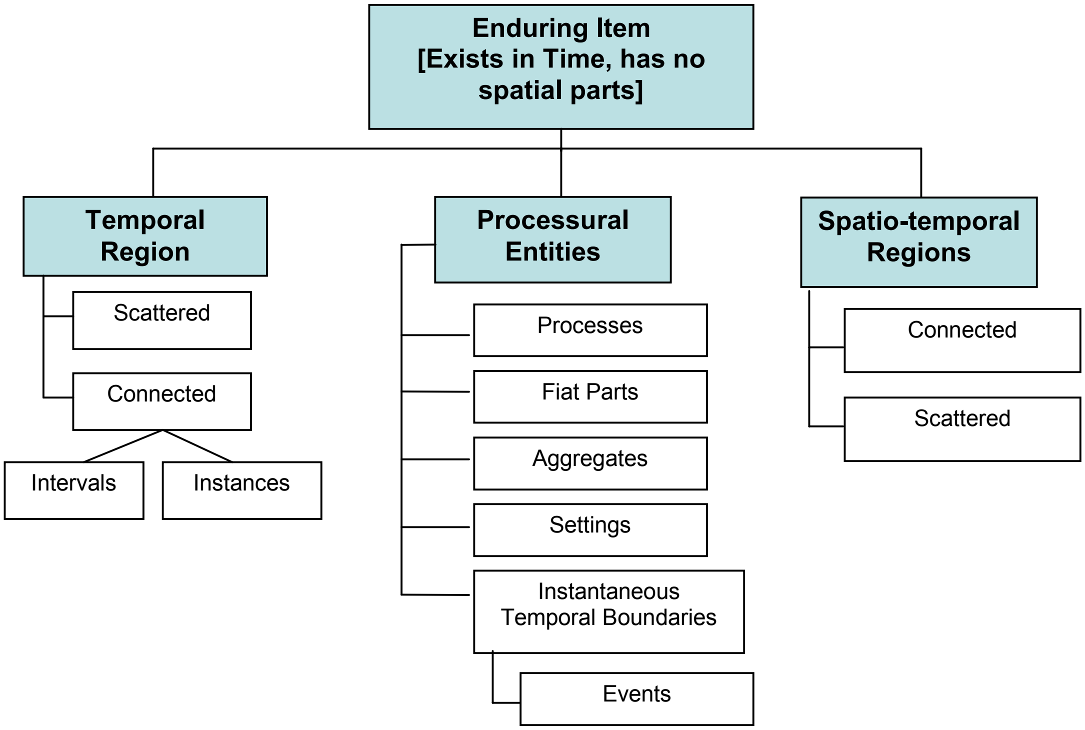

Figure 5. *SPAN BFO upper-level categories*

The boundaries between SNAP objects are normally bona fide (i.e., are independently existing portions of reality). Grenon and Smith point out that, "you are a SNAP entity; your life or history is a SPAN entity. You are three-dimensional; your life is four-dimensional. Moreover, no SNAP entity is ever a part of any SPAN entity and vice versa" (Grenon & Smith, 2003, p. 4).

A comprehensive and accurate understanding of threat can be achieved through the construction of a formal Threat Ontology (ThrO). Consequently, a threat ontology will draw from both SNAP and SPAN, since as pointed out by Roy, Paradis & Allouche (2002), threats are composed of the following three basic components: 1) an intentional component (which in turn is informed by an adversary's interests/desires, capabilities, vulnerabilities, and opportunities); 2) a plan component (containing a method and sub-goals), the sum of which could result in; 3) performed actions and consequences (both of which are supported by past and future activities of the adversary). The first and second threat components contain mostly SNAP entities, which share relations with certain SPAN counterparts (e.g., any intention contains both an intentional object (that which is intended – be it mental or physical) and an intentional act (the cognitive manner in which it is intended)). The third component, conversely, contains primarily SPAN components, which share relations with certain SNAP items (since any action occurs within a defined spatial region, performed by some agent or group of agents, etc.). A threat ontology must, therefore, draw from certain transcategorical ontological relations which bridge the gap between SNAP and SPAN (i.e., the relations between continuants and occurrents). Several such transcategorical formal relations have been described by Smith and Grenon (2003).

4. Threat ontology

Information fusion is divided into 5 sub-levels of investigation (Levels 0-4). Level 0 concerns sub-object data assessment (signal-object-observable states based on pixel/signal level data association and characterization). Level 1 concerns object assessment (estimation and prediction of entity states based on observations). Level 2 concerns situation assessment (estimation and prediction of entity states based on inferred entity relations). Level 3 concerns impact (threat) assessment (estimation and prediction of effects on situations based on plans and activities performed by all participants – friendly or non-friendly). Level 4 concerns process refinement (adaptive data acquisition and processing to support mission objectives). Threats are objects of agent-based complex activities, which result in some impact on other agents, objects, political structures, financial institutions, etc., and are the focus of level 3 fusion. However, in order to understand items at level 3, it is important to understand the items at Levels 0-2, since signal quality (L0), objects (both physical and nonphysical) (L1), and object attributes and relations (L2) all play an important role in reasoning about Level 3 items (those processes aimed at perceiving and representing other agents, their purposeful activities, impacts, or threats).

The United States Department of Defense (US DoD) defines a viable (i.e., actual) threat as being composed of 1) an intent, 2) a capability (i.e., lethality), and 3) an opportunity. A threat, therefore, is a relational item, meaning that its existence is tied to the existence of other items standing to one another in various kinds of ways. For example, if one were to remove or alter the three relational elements of a threat, the threat, as a whole, would be altered or destroyed. Thus, the existence of threats relies on the existence of their parts or members and the relations that obtain between them.

4.1 The concept of an integral whole

Since level 3 items such as threats contain both SNAP and SPAN constituents, as well as relations, it is unlikely that a complete ontological description of Level 3 items can be derived solely by examining and modeling aggregates, because aggregates do not have an intrinsic relational structure to their parts. Aggregates are merely collocations of independent items, whereas wholes, such as threats, hang together through dependence relations. This means that wholes have dependent parts, whereas aggregates, collections, sums, or lists of items do not. In this sense, tank battalions are not merely aggregates of tanks, human beings are not simply aggregates of cells, terrorist organizations are not merely aggregates of non-friendly, unconventional adversaries, and threats are not merely aggregates of intentions, activities, agents, and weapons of various sorts. Tank battalions, human beings, terrorist organizations, and especially threats are wholes that possess parts whose organization cannot be exhaustively explained by the concepts of summation or aggregation, since they are bound together through strict laws of dependence. Wholes of this kind are called *integral wholes*.

The concept of an integral whole can be traced to Edmund Husserl's *Logical Investigations*, which is the first earnest treatise on formal ontology (1900-01). It has since been expanded by Peter Simons (1987) and others into a robust, axiomatized, mereology capable of being expressed in first order logic. An integral whole is one whose part-organization is itself an essential part of the whole. An essential part is one whose existence is necessary for the whole's existence (a non-essential part is one whose existence is contingent). A viable threat is therefore an integral whole, which possesses three essential and interdependent parts: its intent, its capability and its opportunity.

Simons explains that for an item to count as an integral whole, it must meet the following four conditions:

1. It must consist of several parts.

2. These parts must stand to one another in certain formal relations.

3. These relations must connect the parts to one another.

4. The total relations of all the parts are characteristic for the kind of complex in question (i.e., one unifying attribute) (Simons, 1987, pp. 354-55).

Condition one rules out purely simple (i.e., theoretically atomic) items, such as Cartesian simples (i.e., purely atomistic concepts), since they contain no constituent parts. Condition two rules out the possibility of treating structured, integrated wholes as mere aggregates, which do not possess intrinsic part-relations. Condition three rules out 'disconnected sums,' illegitimate wholes such as my thumb and your foot, where there is no clear or obvious relationship between the parts of that whole. Condition four refers to the Gestalt (i.e., shape, figure, or form) of the whole in question and addresses the idea that a structured (integral) whole must possess one essential, overarching characteristic that is founded on the harmony of its parts (Little, 2002).

Threats form a special kind of tri-partite integral whole (see Figure 6). Its constituent parts form necessary dependence relations to one another, such that if one of the parts is removed, the entire whole is dissolved. This means that the parts are inextricably related on one another by a special type of ontological dependence – that being foundational dependence. Foundational dependence says that if an A is founded on a B, then A cannot exist without B existing, since A's very existence is founded on the existence of something else (B).

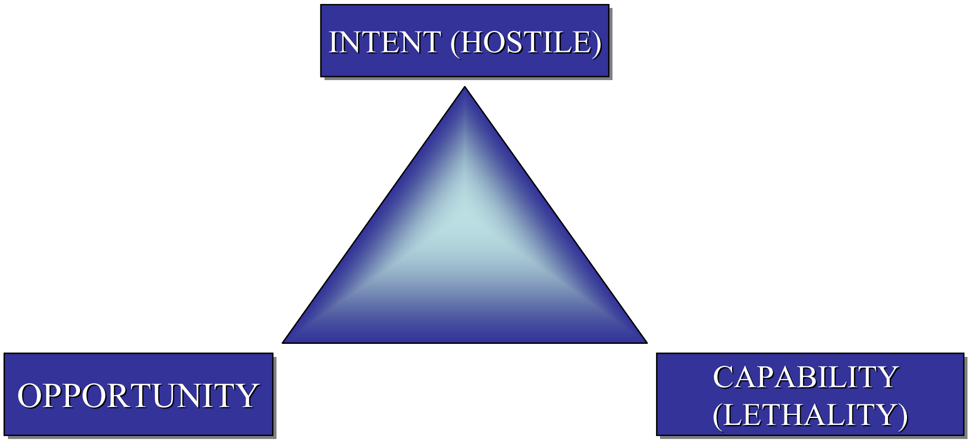

Figure 6. *The tri-partite model of threat*

Since threats contain parts which are tied together through foundational dependence relations, we can treat each element of the threat as a *center of gravity*, meaning that each element is a critical component of the threat condition, disruption of which can result in removing the will, the capacity, or the ability to carry out the threat as planned.

4.2 Decomposing threat components – the lexical task

The first step in providing an ontology of threat is to provide a sufficiently detailed lexicon of terms, which can subsequently be used for constructing an initial taxonomy. Taxonomies should not be understood as full-blown ontologies, since the former can be rather loosely, informally structured clusters of data, whereas the latter come with formal relations (usually of subsumption, dependence, inheritance, etc.) built in. By examining legitimate lexical sources, one can often provide a very complete, though sometimes conflicting, set of definitions from which to begin the ontology-engineering process. Since the U.S. DoD understanding of threat involves three complex terms, it is often useful to decompose these terms further in order to flesh out certain implicit definitions which may be of interest.

For example, by decomposing the terms *capability*, *opportunity*, and *intent*, one can arrive at more nuclear terms which depict certain related groups of occurrents and continuants, which in turn contain both SNAP and SPAN constituents. Consider, for example, the Oxford English Dictionary's (OED) definition of 'capability.' The OED defines a 'capability' as follows:

Capability (df.)=

1. The quality of having room for any thing; ability to receive or contain. CAPACITY.

2. Power or ability in general, whether physical or mental; capacity.

3. Legal or moral qualification or capacity.
4. The quality of being susceptible of, or admitting of treatment, in any specified manner.
5. (usually pl.) An undeveloped faculty or property; a condition, physical or otherwise, capable of being converted or turned to use.

Notice that the term *capability* is synonymous (or at least closely related to) the term *capacity*. This allows for a further decomposition of the semantic value of capability to include capacities as well. The OED defines *capacity* as:

Capacity (df.)=

1. gen. The power, ability, or faculty for anything in particular.
2. The quality or condition of admitting or being open to action or treatment; capability, possibility.

By combining the definitions of capability with those of capacity, we are able to tease out a group of SNAP items (e.g., physical power, mental power, legal qualification, moral qualification, qualities of susceptibility or treatment, physical faculties and conditions, mental faculties and conditions, conversion potential, etc.) as well as a group of SPAN items (e.g., powers, abilities, faculties, actions, treatments, taken as temporally unfolding items). By combining the SNAP and SPAN items, one can then transform this informal set of definitions into a robust, formally structured model capable of depicting numerous types of relations necessary for a clearer metaphysical understanding of capability.

The same activity can then be performed on the other two terms – intent and opportunity, resulting in a formal ontological description of those elements, which comprise a threat in its most basic form. Moreover, by working in even larger definitions of threat, one can engineer a formal framework, which covers numerous layers of granularity (e.g., socio-cultural layers, physical layers, logical layers, etc.).

As stated earlier, the purpose of BFO is to allow for numerous veridical perspectives on reality, all of which help to provide an overarching and comprehensive world-view. An example of an extension of the concept of threat, one aimed at human assessment (i.e., threat analysis), can be found in Roy et al (2002). They define threat analysis as follows:

> *The analysis of the past, present and expected actions of external entities, and their consequences, to identify menacing situations and quantitatively establish the degree of their impact on the mission, the intents, the plans, the actions and the human and material assets of some valuable units to be protected, taking into account the defensive actions that could be performed to reduce, avoid or eliminate the identified menace.*

Here too, we can tease out the various SNAP – SPAN items, adding them to our list above, in turn adding some much needed elements of human analysis such as: 1) impacts on missions, 2) protection of valuable assets, 3) performance of defensive actions, and 4) essential elements of information (EEI's). Investigation of this sort allows for an even more robust metaphysical description of the world, one that includes levels of human interaction such as strategy, preemption, prediction, and reaction (Roy et al, 2002).

4.3 Dispersed wholes

The parts or members of threats are often extended over spatial regions or temporal periods. Items such as threats must therefore be understood as dispersed wholes, meaning wholes which are spatially, temporally or causally non-contiguous, but which nonetheless contain parts or members that stand to one another via a certain unifying feature or characteristic of that whole. Examples of dispersed wholes include armies, French speakers, island chains, and the solar system. Dispersed wholes appear to exist as mere aggregates of items, but actually are constituted by much stronger ontological relations. As stated above, aggregates are simply collections of things, similar to sets, which contain members. Because of the loose connectivity of their members, aggregates cannot maintain their identity over time the way that more legitimate wholes can. Wholes, unlike aggregates, possess parts, which stand to one another, and to the whole, as relational items, bound together by their part structure. They can maintain their identity over time, in spite of the fact that they may gain or lose members. As long as they do not lose their essential parts (responsible for making them identifiable as what they are), they survive from moment to moment, even though they can have ever changing constituents. Aggregates, on the other hand, cannot maintain their identity if they gain or lose members, since at each moment, they are only the summation of their specific members at that time and in that place. Likewise, sets, if the constituents of an aggregation change, then by definition, the aggregate changes. Conversely, if contingent parts of a whole change (e.g., an army loses some soldiers in battle, a nation loses members due to death), the whole can remain intact and maintain its identity as the U.S. Army or the country of Canada, in spite of such changes.

The structural integrity of dispersed wholes (their tied-togetherness) can be defined by certain specific (though perhaps contingent) metrics of the whole's characteristics. Different metrics can be defined for different characteristics, based on certain contingencies within the environment, the agent doing the measuring, etc. Thus, even though some whole in question possesses parts, which stand to one another in certain formal part-relations, the proximity metric used (which could measure characteristics, spatial locations, behaviors, causal relations) could vary. In this sense, the proximity metric will be contingently affected by domain-specific content pertaining to the given situation.

By ontologically defining certain essential characteristics of an item in question, then defining a metric and threshold for that metric, is hypothesized that dispersed wholes of various consistencies will be more easily identified and defined. The following methodology is being developed in the ThrO in order to help identify and describe different types of dispersed wholes:

Let $A = \{a_i\}$ be a set of ontologically defined essential attributes of the items of interest E and $P = \{p_i\}$ be a set of metrics (both physically and psychologically grounded) measuring the proximity of those essential attributes. Then a combination of these metrics $D = F(p1, \ldots, pi)$ can provide us with a measure of proximity for a dispersed whole, which will define the level of integrity of the dispersed whole under consideration. Now we need to define a threshold T or a set of thresholds we consider each attribute separately $\{t_i\}$, which will establish a boundary separating this dispersed whole from an aggregate: E is a dispersed whole if $D < T$. The existence of this boundary is guaranteed by Hegel's law for transitioning quantity into quality that says that purely quantitative differences beyond a certain point transform into qualitative changes. Hegel's law of transition from quantity into quality can be used to define special types of relations namely "quantity-quality" relations. It also provides an initial starting point for investigating approaches to designing a computationally tractable ontology useful for various fusion applications, specifically those related to the detection and increased ability to monitor dispersed wholes such as terrorist organizations, dispersed materials used in weapons proliferation, etc. However, much more research needs to be done on this topic.

5. Mereotopology of ThrO

The formal structure of an ontology provides a rigorous metaphysical description of certain qualitative aspects of the world. Fusion can use these kinds of qualitative aspects defined within the ontology for use in developing improved quantitative representations of the ontology's quantitative categories. The ThrO's formal structure has been carefully designed to ensure that the upper-level formal structure will easily lend itself to supporting quantitative metrics in future applications.

The metaphysical, upper-level design of the ontology requires a two-fold approach for its development. First we employ a *mereology*, which is an axiomatized theory of parts, part-relations, contact, separation and boundary. Second, we employ, in tandem with the mereology, a *topology*, which is a formal theory of wholes that provides the capabilities to deal with dispersed spatial items and spatio-temporal extensions. The combination of mereology and topology result in a larger unified theory of *mereotopology*, which is a theory capable of treating the ontological complexities of items such as threats (Smith, 1996; Casati and Varzi, 1999).

5.1 Mereological primitives

ThrO contains the following mereological primitives:

1. is a part of

2. is necessary such that

3. foundation

The first mereological primitive, *is a part of* operates under the common sense belief that items in the world have constituent members, pieces, properties or attributes which compose them. The basic conceptual connection between parts and wholes can be understood in the following manner:

> *x is said to be a whole insofar as there is some y, which is a component part of x, and x, and y are not identical.*

The second mereological primitive, *is necessary such that* implies modal necessity, which allows for the capability of making universal statements that hold for all possible worlds. Necessary features of the world could not be otherwise.

The third mereological primitive of *foundation* is a theory of the relation between a priori connected parts. The concept of foundation can be defined as follows:

> *x is founded on y =df. x can only exist if y exists.*

Foundation relations are used to show certain put-togethered features of the world where one item's very existence relies upon (i.e., is founded upon) some other item's existence.

5.1.1 Formal properties of the part-relation

The part-relation relates two or more *individuals* (Leonard and Goodman, 1940). Individuals are distinct items that "are all of the lowest logical type, as distinct from entities of higher types such as classes, functions or attributes" (Simons, 1987, p. 10). Thus any two individuals, regardless of type or kind, can stand in a part-relation. For example, every living organ is a part of some specific organism, every citizen is a part of some nation, every atom is a part of some substance, and most generally, every object is a part of reality.

There are three distinct formal properties that hold between individuals in a part-relation. These formal properties constitute the meaning of the word 'part' in such a way that "anyone who seriously disagrees with them has failed to understand the word" (Simons, 1987, p. 11). The following three axioms provide the formal properties of a part-relation:

> A1. Asymmetry. If x is a proper part of y, then y is not a part of x.
>
> A2. Transitivity. If x is a proper part of y, and y is a part of z, then x is a part of z.
>
> A3. Irreflexivity. Nothing is a proper part of itself.

The axiom of asymmetry (A1) provides the basis for a theory within which we can formulate different kinds of nesting. A1 reflects reality, in that if some single terrorist (x) is part of some terrorist network (y), then that network (y) cannot be a part of that terrorist (x).

The axiom of transitivity (A2) provides the formal basis for representing two parts whose relation is mediated by a third part. For example, if some threatening intention (x) is part of some person's belief structure (y), and that person's belief structure is part of a terrorist doctrine (z), then that intention (x) is also part of that person's terrorist doctrine (z).

The axiom of irreflexivity (A3) simply states that parthood is not a reflexive property. A whole cannot be understood as a part of itself. For example, a person's body is not a part of itself. It simply is itself. If a whole could be part of itself, then the whole and one of its parts would be identical, which is absurd. In these cases, the attribute of being a part and the attribute of being a whole could be attributed to the same item at the same time.

5.1.2 Extended mereological concepts

The following mereological concepts are derived from the basic concept of part relation. If we examine our perceptions of everyday items, we find that the following mereological concepts are entirely commonplace. The world is directly perceived as a place in which items exhibit various sorts of spatial and temporal relations to one another. For example, when animals perceive an occluded object (i.e. an object which is partly hidden from view), they immediately recognize that their perception is of two discrete objects standing in a spatial relationship, such that one is in front of the other. One object occupies the perceptual foreground and the other occupies the perceptual background. The foreground object appears to overlap the background object, blocking a portion of it from perception. Basic mereological relations such as these are so commonplace that they can be cognized by infants and many species of animals.

The following list of mereological concepts is not exhaustive, rather it represents those concepts that are essential to designing a proper threat ontology.

- Overlap. (x o y) x overlaps y =df. there is some z such that z is a part of x and z is a part of y.

- Discreteness. (x | y) x is discrete from y =df. there is no part of x that overlaps any part of y.

- Binary Product. (x · y) the product of x and y =df. that individual which is part of both x and y, and which is such that any common part of both x and y is part of it.

- Mereological Sum (x + y) the mereological sum of two individuals, x and y, =df. that individual which something overlaps if, and only if, it overlaps at least one of x and y.

- Mereological Difference (x – y) if x and y are two individuals, then their mereological difference =df. the largest individual contained in x which has no part in common with y.

- The Universe. (U) the universe =df. the mereological sum of all objects.

- Complement. (U–x) the complement of x =df. the sum of all objects which are discrete from x.

5.2 Topological primitives

We must provide a theory that is capable of distinguishing between wholes whose parts are integrated and wholes whose parts are scattered (but which can nevertheless be conceived of as unitary items). Wholes whose parts are scattered are dispersed wholes and it is these types of wholes, which require the supplementation of mereology with a theory of topology.

The topological primitives used in ThrO are:

- Connectedness.
- Spatial extendedness.

Connectedness can be defined in terms of boundary, since it, like the concept 'is a part of' refers to a pre-theoretical, common-sense understanding of the world and the way the world is put together. Connectedness is both reflexive and symmetric, but not transitive. Connectedness is reflexive, since every object is connected to itself. Connectedness is symmetric, in that, if x is connected to y, then y is also connected to x. Connectedness is non-transitive, in that, if x is connected to y, and y is connected to z, it does not follow that x is connected to z. For example, Canada is connected to the United States, and the United States is connected to Mexico, but Canada is not connected to Mexico.

Spatial extendedness is another topological primitive, since our perceptual systems are designed to provide us with a common sense understanding of a spatially extended world, which serves as the backdrop for all of our experiences. Neuronal plasticity, for example, shows that were it not for the inherent structure of our surroundings, our perceptual systems would not have developed the way they, in fact, have.

5.3 Combining topology and mereology

The result of supplementing mereology with topology is *mereotopology*. However, simply adding the ground topological concept of connectedness to the mereological axioms of A1-A3 above does not yield a mereotopology, unless one adds a bridging principle to link them. The bridging principle suggested by Casati and Varzi (1999), is *monotonicity*. Monotonicity can be defined as follows:

> *Monotonicity =df. If some thing x is part of some thing y, then whatever is connected to x is connected to y.*

According to the definition of monotonicity, parthood is a sufficient condition for connection, but connection is not a sufficient condition for parthood. If the head is part of the body, then they must be connected. However, if a hat is connected to someone's head while they are wearing it, the hat does not become part of the head or the body.

The concept of monotonicity also implies that mereological overlap is a form of connection. Thus we may say that:

> *If x overlaps y, then x is connected to y.*

Overlap, therefore, entails connectedness. For example, if your left hand overlaps your right hand, then the two are spatially connected. Conversely, connection does not entail overlap. For example, Germany and Denmark are connected to one another, yet Germany and Denmark do not overlap.

We now have a defining condition for an external topological connection, where x shares none of its parts with y. We may define an external connection as:

> *x and y are externally connected =df. x is connected to y and x does not overlap any part of y.*

External connection is used to define those relations where two things abut, barely touch or come into contact with one another, but where no parts of either object overlap. Casati and Varzi point out that the relation of external connection is symmetric, but is neither transitive nor reflexive. For example, black and white squares on a checkerboard are externally connected, since they abut one another, but no part of any white square overlaps any parts of any black square. External connection is a symmetric relation in that, if a black square is externally connected to a white square then the white square is externally connected to the black one. External connection is a non-transitive relation in that, if two black squares are externally connected to one white square, which is between them, the black squares are not externally connected to each other. Finally, external connection is non-reflexive in that, nothing can be externally connected to itself.

6. Using ThrO for domain-specific applications

In this chapter, we introduce different types of threats, distinguish between conventional threats and unconventional threats, and illustrate the theoretical concepts with a case study.

6.1 Decomposition of threat types

Threats can exist in various forms and be composed of various types of specific elements. However, the primary goal of ThrO in its initial stages of development is to ask, what threats are metaphysically, or in other words, what are the qualities common to all threats in specie? In order to answer this question, it is important to understand the formal structure of threats themselves, regardless of their domain-specific characteristics. This allows us to have confidence that our analysis of threat is not too narrow in scope, or that our model is not based on overly specific states of affairs, which only apply to limited threat scenarios.

6.1.1 Potential and viable threats

Threats can first be decomposed into potential and viable threats (See Figure 7).

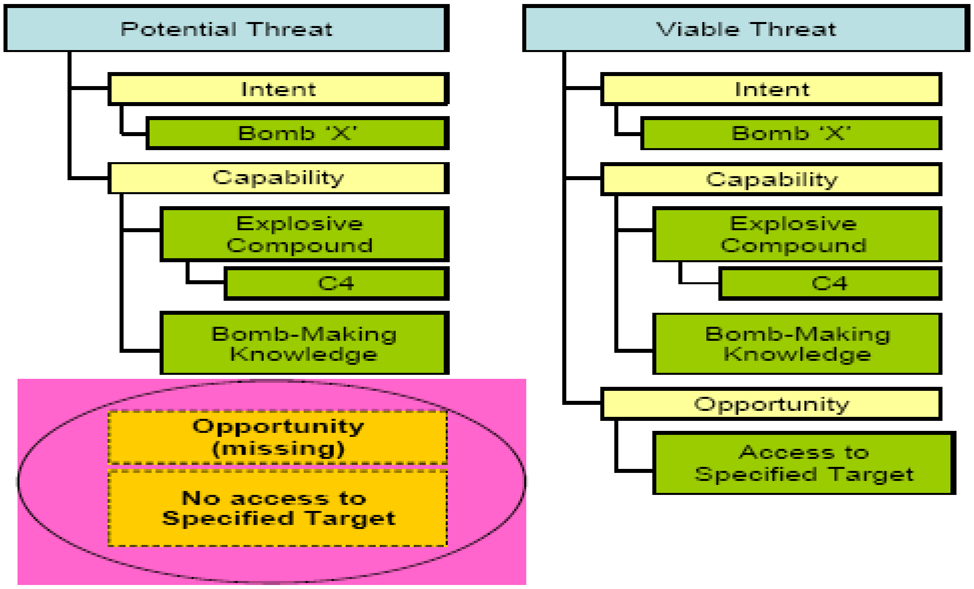

Figure 7. *Potential threat vs. viable threat*

Potential threats are threats that are missing an essential part (either the intent, capability, or opportunity). Potential threats are threats, which are not in a state of being, but are only in a state of becoming, meaning they are not yet actualized. A viable threat, on the other hand, is the type of threat that has been discussed throughout this report. A viable threat contains all three essential elements of intent, capability, and opportunity and all foundational relations between them. Viable threats contain both SNAP and SPAN components, whereas potential threats are SPAN items only, because their existence is only temporally, not physically, instantiated.

6.1.2 Conventional vs unconventional threats

We have decomposed both potential and viable threats into the categories of conventional and unconventional. Conventional threats apply to typical warfare scenarios between nation-states which agree to certain battlefield codes of conduct that rule out certain types of activities, weapons, and protocols which have been agreed upon by recognized governing bodies (U.N., NATO) as illegitimate forms of conflict. Unconventional threats are threats, which do not adhere to any set of pre-established rules of warfare. Unconventional threats can include: intentional injury to non-military personnel (e.g., children, civilians), intentionally inflicted damage to non-military targets (buildings, financial institutions, computer networks), inhumane treatment to military and nonmilitary personnel (e.g., torture), or dispersal of unconventional weapons (e.g., radioactive bombs, biological agents, chemical agents).

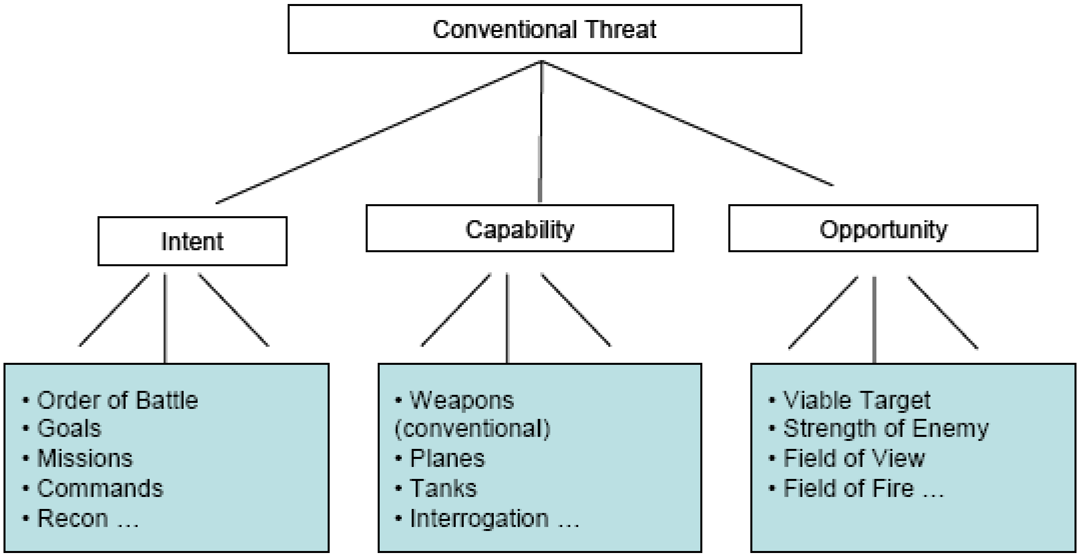

Figure 8. *Basic formal ontology for conventional threat*

All threats, conventional or unconventional, have the same formal structure (i.e., the same tri-partite whole structure including intent, capability, and opportunity) at the upper-levels of abstraction (see Figures 8 and 9). Their differences are discovered in their domain-specific instance types.

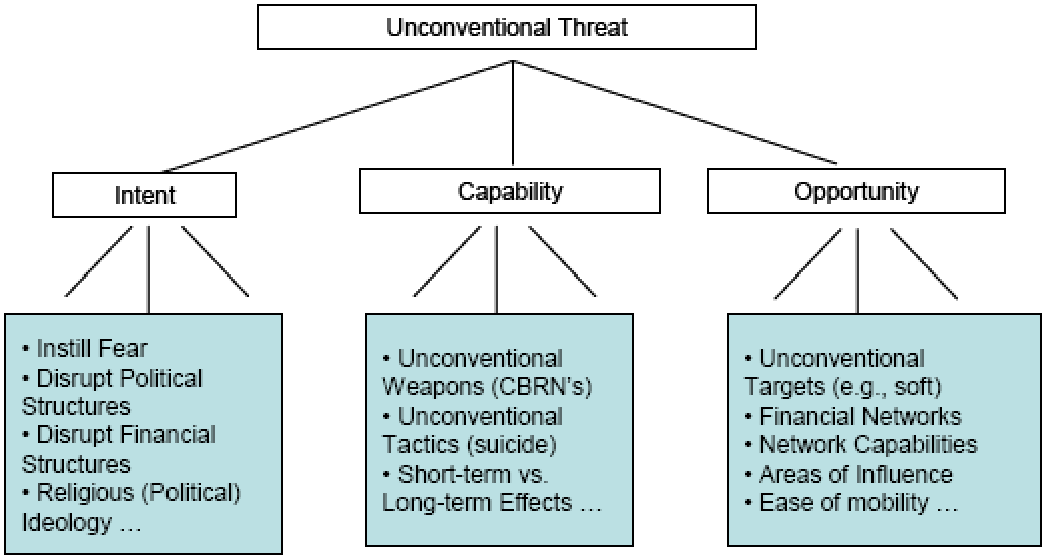
Unconventional Threat
Intent
Capability
Opportunity
• Instill Fear
• Disrupt Political Structures
• Disrupt Financial Structures
• Religious (Political) Ideology …
• Unconventional Weapons (CBRN's)
• Unconventional Tactics (suicide)
• Short-term vs. Long-term Effects …
• Unconventional Targets (e.g., soft)
• Financial Networks
• Network Capabilities
• Areas of Influence
• Ease of mobility …

Figure 9. *Basic formal ontology for unconventional threat*

This means that there will be different lexicons needed for the lower ontology levels for various types of threats, since different application-specific domains will include different instances of agents, weapons, targets, activities, etc. Likewise there will be potentially different values assigned to the different lexical items at the ontology's lower levels. The modeling of the lower levels and the values assigned for fusion purposes will be informed by expert knowledge for all such kinds of specific threats identified.

6.2 Dirty bomb case study

As a test bed for ThrO, we have designed a case study aimed at a particular form of unconventional threat – a dirty bomb. A dirty bomb is a conventional explosive (dynamite, C4, etc.), which has been laced with a radioactive element. While the effects of a dirty bomb attack are, as of yet, not entirely known, it is assumed that there would be negative short-term effects, including: 1) personal injuries from the blast itself, 2) ensuing panic (especially in large metropolitan areas), 3) possible building or infrastructural damages. The long-term effects could be more far-reaching and may include: 1) personal injury due to radiation sickness (both short-term and long-term) or cancer 2) extensive HAZMAT clean-up costs, 3) decrease in property value due to perceived unsafe living conditions, 4) prolonged psychological effects due to perceived unsafe living conditions (www.fas.org).

The following list represents interesting features of a dirty bomb, which may make this a candidate weapon for terrorist groups or parties interested in carrying out unconventional types of threats against various nation-states:

- Readily available materials from numerous sources in the private sector including: medical equipment, food irradiation plants, sterilization products, cancer treatment products. (increased opportunity)
- Easy to construct (increased opportunity)
- Easy to conceal (increased opportunity)
- Potential for large-scale physical and psychological injury (increased capability)
- Potential for both short-term and long-term effects (increased capability)
- Potential for high amount of secondary threats (food/water contamination, soil contamination) (increased capability).

6.2.1 The mereology of dirty bombs

The example of a dirty bomb depicts how the three elements of a threat exist to one another as mereological parts of an overarching integral whole. The part-relations between intention, capability and opportunity are such that enhancement of one element, can result in the enhancement of one or more of the others. For example, the availability of materials, coupled with the relative ease of construction and concealment, increase the opportunity for one to build and disperse a dirty bomb. The current high access to both materials and potential targets, results in this being an increasingly potential sort of threat. The potential for large-scale physical and psychological injuries as well as various short- and long-term effects amounts to an increased capability (i.e., lethality) of these types of weapons. The potential to inflict damage at various levels, to various populations, over extended time periods, and with increasing chances of secondary threats, also results in this being an increasingly viable sort of threat. The increase in opportunities and capabilities could lead to a direct increase in intentions to use such weapons, since the foundational relations between intentions, capabilities, and opportunities are such that an increase in one item can often affect an increase in other interrelated elements of a whole. Such reciprocal increases in intentions, capabilities and opportunities point to the need for a solid foundational theory of mereology in order to model such relations within the ontology.

6.2.2 The mereotopology of dirty bombs

Because the possible elements of a dirty bomb are: 1) numerous, 2) found in an array of industrial and medical sources, and 3) seemingly spatially (and causally) unrelated items until assembled together, such elements can be seen as representing a type of dispersed whole. The parts of a dirty bomb may be gained from different places, through different political channels, transported in a disassembled manner, and include different agents in their proliferation, concealment, and construction. Thus arises the need for an ontological model which can model a complete array of various places, times, agents, etc., which may, on the surface, appear to have little or no connection. An ontology of dispersed parts of a dirty bomb could hopefully uncover interesting overlaps or connections between parts of the threat's integral whole (kinds of materials needed to manufacture the bomb, places where these materials are stored, agents may have access to them, places where such materials could be transported, opportunities that exist for their dispersal, stated intentions to use such items, etc.) (See Figure 10). Since information about such disparate items is not currently warehoused in any single frame of reference, it is our hope that a full-blown SNAP/SPAN formal ontology of dirty bombs (and their materials, activities and groups associated with their use) could help for knowledge discovery as to new kinds of potential and viable threats, as well as some prediction of the potential hazards associated with their dispersal.

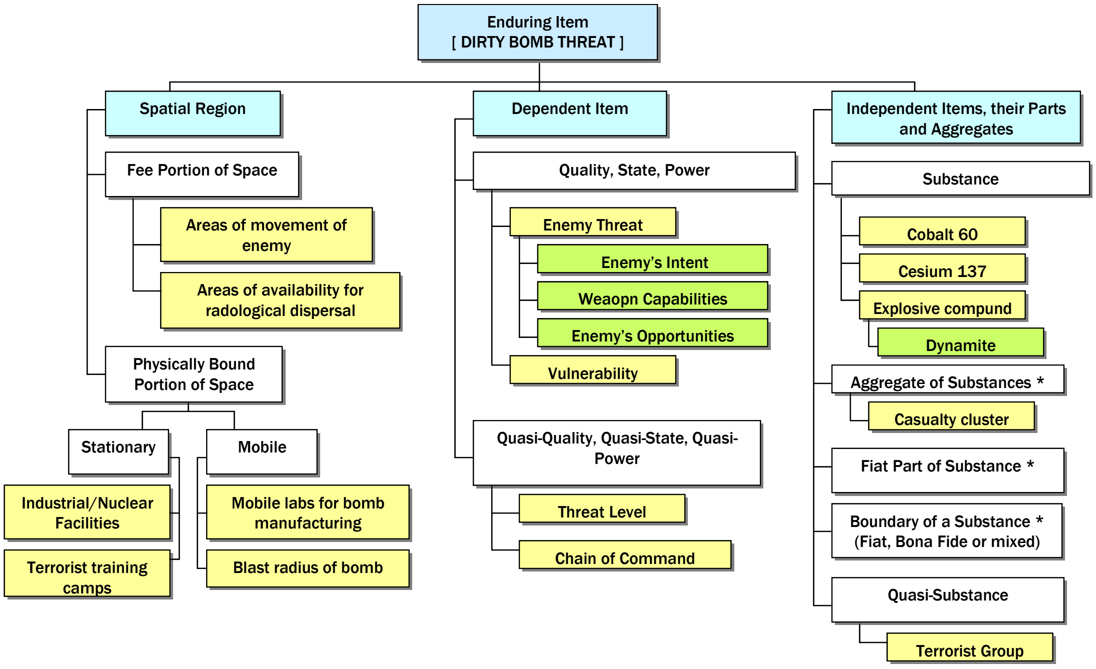

Figure 10. *Rudimentary elements of a SNAP ThrO for dirty bombs*

By applying a mereotopologically grounded ontology to dirty bomb cases, we could attempt to provide an exhaustive list of items, which are prime candidates for inclusion in a dirty bomb threat. For example, to begin, one could construct an exhaustive representational model of all such radiological elements such as Cobalt-60 and Cesium-137, as well as conventional explosives such as TNT or C4, coupled with spatial regions pointing out areas for manufacture and dispersal of the weapon, plume effects, etc (Kelly 2002). Then, by applying some metrics to the ontology, in terms of quantifying certain distances, overlaps, contacts, and connections between these categorical items (understood in terms of certain a priori predetermined thresholds), it is hoped that certain threat values could then be calculated based on evidence gained from the ontology. However, this process is only still in its initial research stages in terms of methodology for the development of such systems. The goal of such an endeavor is to be able to produce improved methods for entire eased identifying, and ultimately neutralizing, threats such as those posed by terrorists with dirty bombs. It had been stated earlier that a decrease or nullification of one of the elements of a threat (as an integral whole) spells the decrease or nullification of the entire viable threat (or in the least, it transforms the viable threat into a potential threat). By analyzing the metrics assigned to the categories within the ontology's branches of capability and opportunity, it is hypothesized here that one could predict, with increased accuracy, an intentional metric, which is directly, associated with these other items. This could prove to be invaluable to thwarting terrorist threats while still in potential. While it is nearly impossible to give quantitative values to qualitative psychological states, it is possible to assign values to capabilities and opportunities afforded to individuals based on behaviorist methodologies. Since the elements of opportunity and capability stand infoundational ontological relations to intentions, one could better infer or predict certain types of adversarial intentions by a form of reverse engineering whereby one determines, within some limited scope, the value of the adversary's intent based on the known values of their capabilities and opportunities.

7. Conclusions

It has been argued here that fusion system design would be enhanced by the development of metaphysically grounded ontologies, which can produce accurate and comprehensive qualitative states about the world. Such qualitative states of the world can be formalized into logically structured frameworks, which provide a basis for transformation into quantitative representations of those states, ultimately providing a set of metrics from which fusion can develop improved algorithms about various states of affairs in the world. The Threat Ontology (ThrO) constructed for this task has been designed as an axiomatized mereotopological item, which is capable of treating parts and wholes of various types, especially difficult types of dispersed wholes whose members are spread out over space and time. In order to test some of the conceptual ontological apparatuses being designed (specifically the theory of threat as a tri-partite whole), a rudimentary dirty bomb case study was analyzed as an initial way to test the connections between intents, capabilities, and opportunities. The connections between these items show, in theory, how an understanding of threat as a whole, containing its three interdependent parts, can be used to increase knowledge about and improve prediction of dirty bomb threats.

Future research is needed to further develop the case study and to fuse the ontology's qualitative categories of existence with quantitative values, providing computationally tractable fusion ontology capable of treating various kinds of threats in various kinds of domains.

References

1. Abidi, M. and Gonzalez, R. (eds.) (1992). Data Fusion in Robotics and Machine intelligence, Academic Press, Inc.

2. Ash, R. (1970) Basic Probability Theory, New York, Wiley.

3. Bedworth, M. and O'Brien, J. (2000) “The Omnibus Model: A New Model of Data Fusion?”, IEEE Aerospace and Electronic Systems Magazine, Volume 15, Issue 4.

4. Benediktsson J, Kanellopoulos, I. (1999) “Classification of Multisource and Hyperspectral Data Based on Decision Fusion,” IEEE Trans. Geosci. Remote Sensing, 36(3), 283-293.

5. Blasch, E. and Plano, S. (2003) “Level 5: User Refinement to aid the Fusion Process,” in Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications, B. Dasarathy (Ed.), Proc of the SPIE, Vol. 5099.

6. Casati, R. & Varzi, A. (1999) Parts and Places, Cambridge, MA: MIT Press.

7. Dennett, D. (1991) Consciousness Explained, Boston: Little Brown and Co.

8. Dasarathy, B. (1997) “Sensor Fusion Potential Exploitation-Innovative Architectures and Illustrative Applications,” IEEE Proceedings, Vol. 85, No.1.

9. Dubois D. and Prade H. (1988) Possibility Theory: An Approach to Computerized Processing of Uncertainty, Plenum Press, New York.

10. Endsley, M. (1995) “Toward a Theory of Situation Awareness in Dynamic Systems” Human Factors Journal., 37, 1.

11. Grenon, P. (2003a) “Knowledge Management from the Ontological Standpoint,” in Proceedings of WM 2003 Workshop on Knowledge Management and Philosophy, April 2003, Luzern Switzerland.

12. Grenon, P. (2003b) “Spatiotemporality in Basic Formal Ontology: SNAP and SPAN, Upper-Level Ontology and Framework for Formalization, IFOMIS Technical Report Series, (http://ifomis.de).

13. Grenon, P. & Smith, B. (2003) “SNAP and SPAN: Towards Dynamic Spatial Ontology,” Spatial Cognition and Computation, 4(1), pp. 69-104.

14. Gruber, T. (1995) “Toward principles for the design of ontologies used for knowledge sharing,” in N. Guarino and R. Poli (eds.), Formal Ontology in Conceptual Analysis and Knowledge Representation.

15.

16. Guarino, N. (1998) “Formal ontology in information systems,” in N. Guarino (ed.), Proc. Of Formal Ontology in Information Systems, pp. 3-15, Trento, Italy, 6-8 June, Amsterdam: IOS Press.

17. Hall, D & Llinas, J. (eds.). (2001) Handbook of Multisensor Data Fusion, Boca Raton: CRC Press.

18. Halpern (2003) Reasoning about Uncertainty, MIT Press.

19. Husserl, E. (1900-01) Logische Untersuchungen, 2 Bde, Husserliana, Band XIX, Den Haag: Martinus Nijoff, 1985 ed.

20. Kelly, H. (2002) “Dirty Bombs: Response to a Threat,” Public Interest Report, Journal of the Federation of American Scientists, Vol. 55, No. 2, March/April.

21. Kokar, M. & Wang, J. (2002) “An example of using ontologies and symbolic information in automatic target recognition,” in Sensor Fusion: Architectures, Algorithms, and Applications VI, pp. 40-50, SPIE.

22. Lambert, D. (2003) Grand Challenges of Information Fusion. Proceedings of the Sixth International Conference on Information Fusion, pages 570-574, Cairns, Australia, pages 213-220, 8 July–10 July.

23. Lenat, D.B. and Guha, R.V. (1990) Building Large Knowledge-Based Systems, Reading, MA: Addison-Wesley Pub. Inc.

24. Leonard, H.S., Goodman N., 1940, "The calculus of individuals and its uses," *Journal of Symbolic Logic 5*: 45-55.

25. Little, E. (2002) Moderate Materialism: Toward a Unified Ontology of Consciousness, Dissertation, State Univ. at Buffalo, Dept. of Philosophy, Center for Cognitive Science.

26. Little, E. (2003) “A Proposed Methodology for Application-Based Formal Ontologies”, Proceedings of the Workshop on Reference Ontologies vs. Application Ontologies, 15-18 Sept., University of Hamburg, CEUR-WS.org.

27. Llinas, J. & Waltz, E. (1990) Multisensor Data Fusion, Boston: Artech House.

28. Menzel, C. & Mayer, R. (1990) IDEF 5 Ontology Description Capture Method: Concept Paper, Research Institute for Computing and Information Systems, Univ. of Houston-Clear Lake.

29. Noy, N. and McGuinness, D. (2001) “Ontology Development 101: A Guide to Creating Your First Ontology,” Stanford University: http://protege.stanford.edu/publications/ontology_development/ontology101-noy-mcguinness.html

30. Ören, T. (2001) “Impact of Data on Simulation: From Early Practices to Federated and Agent-Directed Simulation” in: A. Heemink et al. (eds.) Proc. of EUROSIM 2001, June 26-29, Delft, the Netherlands.

31. Rogova, G., Stomper P. (2002) “Information Fusion approach to microcalcification characterization,” International Journal on Multi-Sensor, Multi-Source Information Fusion, 3 (2), 91-102.

32. Rogova, G. Llinas, J. (1996) “Data Fusion For Real-Time Dynamic Traffic Assignment,” In Proc. of the Sixth Annual IEEE Dual-Use Technologies and Applications Conference, 1-5.

33. Rogova (1994) “Combining the Results of Several Neural Network Classifiers,” Neural Networks, **7**(5), 777-781.

34. Roy, J. (2001) "From Data Fusion to Situation Analysis", Proceedings of the Fourth International Conference on Information Fusion (FUSION 2001), Montreal, Canada, August 7-10.32.

35. Roy, J., Paradis, S., Allouche, M. (2002) Threat Evaluation for Impact Assessment In Situation Analysis Systems, SPIE Proceedings, Vol. 4729, Signal Processing, Sensor Fusion, and Target Recognition XI, Ivan Kadar Ed., Orlando, 1-5 April 2002, pages 329-341.

36. Salerno, J. (2002) “Information Fusion: A High-Level Architecture Overview,” Proc. Intl Conf on Information Fusion (FUSION2002), Annapolis, Md.

37. Searle, J. (1992) The Rediscovery of the Mind, Cambridge, MA: The MIT Press

38. Scott, P. and Rogova, G. (2004) “Crisis management in a Data Fusion Synthetic Task Environment,” in: Proc. of FUSION 2004, Stockholm, Sweden.

39. Shafer, G. (1976) A Mathematical theory of evidence, Princeton, NJ.

40. Shimshoni Y. and Intrator N. (1998) “Classification of Seismic Signals by Integrating Ensembles of Neural networks,” IEEE Trans. on Signal Processing, 46(5), 1194-1201.

41. Simons, P. (1987) Parts: A Study in Ontology, Oxford: Oxford Univ. Press.

42. Smets, P. and Kennes, P. (1994) “The transferable belief model,” Artificial Intelligence, 66, 191-243.

43. Smith, B. (1989) “Logic and Formal Ontology,” in J. N. Mohanty and W. McKenna (eds.), Husserl’s Phenomenology: A Textbook. Washington, D. C.: Center for Advanced Research in Phenomenology and University Press of America, 29-67.

44. Smith, B. (1996) “Mereotopology: A Theory of Parts and Boundaries,” Data and Knowledge Engineering, 20 (1996), 287–303.

45. Smith, B., 1997, “Boundaries: An Essay in Mereotopology”, in L. H. Hahn (ed.), *The Philosophy of Roderick Chisholm*, Chicago and La Salle, IL: Open Court, pp. 534–61.

46. Steinberg, A.N., Bowman, C.L. and White, F.E. (1998), “Revision to the JDL data fusion model”, Joint NATO/IRIS Conference, Quebec City, Canada, October 1998

47. Tangney, J. (2002) AFOSR Programs in Higher Levels of Information Fusion, Fusion 2002, Annapolis MD, July 8-11, pp 557-561. Also in IEEE Aerospace and Electronics Systems Society Magazine, pp. 21-25, Nov 2003.

48. Uschold, M. and Gruninger, M. (1996) “Ontologies: Principles, Methods and Applications,” The Knowledge Engineering Review, 11(2), pp. 93-136.

49. Welty, C. and Smith, B. (eds.) (2001) Formal Ontology and Information Systems, New York: ACM Press.

50. White, F. (1988) “A Model for Data Fusion,” in Proceedings of SPIE 1st National Symposium on Sensor Fusion, Vol. 2, Orlando, Florida.

This page intentionally left blank.

Internal distribution

1	-	Director General
3	-	Document Library
1	-	Head/C2 Decision Support System
1	-	Head/Intelligence and Information
1	-	Head/System of Systems
1	-	A.-C. Boury-Brisset (author)
1	-	M. Allouche
1	-	M. Blanchette
1	-	M. Bélanger
1	-	J. Berger
1	-	A. Boukhtouta
1	-	A. Benaskeur
1	-	R. Breton
1	-	C. Daigle
1	-	E. Dorion
1	-	A. Guitouni
1	-	H. Irandoust
1	-	A.-L. Jousselme
1	-	P. Maupin
1	-	S. Paradis
1	-	F. Rhéaume
1	-	J. Roy
1	-	P. Vallin
1	-	A. Auger
1	-	Ltv. L. St-Pierre
1	-	LCdr E. Woodliffe

External distribution

1 - Library and Archives Canada

395 Wellington Street, Ottawa, ON, K1A 0N4

1 - Director Research and Development Knowledge and Information Management (PDF file)

1 - Director Science and Technology (C4ISR)

Constitution Building, 305 Rideau St., Ottawa, ON,
K1N 9E5

1 - Director Science and Technology (Air)

Constitution Building, 305 Rideau St., Ottawa, ON,
K1N 9E5

1 - Director Science and Technology (Land)

Constitution Building, 305 Rideau St., Ottawa, ON

K1N 9E5

1 - Director Science and Technology (Maritime)

Constitution Building, 305 Rideau St., Ottawa, ON

K1N 9E5

1 - Director Maritime Requirements Sea (DMRS)

Louis St.Laurent Bldg, 555 Boul. de la Carrière,
Gatineau, QC, J8Y 6T5

1 - Director Maritime Requirements Sea 4

Louis St.Laurent Bldg, 555 Boul. de la Carrière,
Gatineau, QC, J8Y 6T5

1 - Director Maritime Requirements Sea (DMRS) 6

Louis St.Laurent Bldg, 555 Boul. de la Carrière,
Gatineau, QC, J8Y 6T5

1 - Director Maritime Requirements Sea (DMRS) 6-2

Louis St.Laurent Bldg, 555 Boul. de la Carrière,
Gatineau, QC, J8Y 6T5

1 - Director Maritime Ship Support (DMSS) 6

Louis St.Laurent Bldg, 555 Boul. de la Carrière,
Gatineau, QC, J8Y 6T5

1 - Director Maritime Ship Support (DMSS) 8

Louis St.Laurent Bldg, 555 Boul. de la Carrière,
Gatineau, QC, J8Y 6T5

2 - Defence R&D Canada – Atlantic:
9 Grove Street, Darmouth, NS, B2Y 3Z7
Attn: Dr. J.S. Kennedy
Dr. B. MacArthur

1 - Canadian Forces Maritime Warfare School
CFB Halifax
PO Box 99000 Stn Forces Halifax,
Nova Scotia, B3K 5X5
Attn: Commanding Officer

1 - E. Little (author)
Calspan-UB Research Center, Inc.
PO Box 400
4455 Genesee Street
Buffalo, NY 14225, USA
Tel.: (716) 829-7841

1 - G. Rogova (author)
Calspan-UB Research Center, Inc.
PO Box 400
4455 Genesee Street
Buffalo, NY 14225, USA
Tel.: (585) 624-1364

This page intentionally left blank.

DOCUMENT CONTROL DATA

1. ORIGINATOR (name and address)
Defence R&D Canada Valcartier
2459 Pie-XI Blvd. North
Québec, QC
G3J 1X8

2. SECURITY CLASSIFICATION
(Including special warning terms if applicable)
Unclassified

3. TITLE (Its classification should be indicated by the appropriate abbreviation (S, C, R or U)
Theoretical foundations and proposed applications of Threat Ontology to information fusion

4. AUTHORS (Last name, first name, middle initial. If military, show rank, e.g. Doe, Maj. John E.)
Little Eric, Rogova Galya, Boury-Brisset Anne-Claire

5. DATE OF PUBLICATION (month and year)
2008

6a. NO. OF PAGES
33

6b .NO. OF REFERENCES
50

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. Give the inclusive dates when a specific reporting period is covered.)
Technical Report

8. SPONSORING ACTIVITY (name and address)

9a. PROJECT OR GRANT NO. (Please specify whether project or grant)

9b. CONTRACT NO.

10a. ORIGINATOR'S DOCUMENT NUMBER
TR 2005-269

10b. OTHER DOCUMENT NOS
N/A

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

☒ Unlimited distribution
□ Restricted to contractors in approved countries (specify)
□ Restricted to Canadian contractors (with need-to-know)
□ Restricted to Government (with need-to-know)
□ Restricted to Defense departments
□ Others

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.)
Unlimited

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

Information Fusion is the process of utilizing information from various sources to produce knowledge about objects and situations in order to assist humans in decision making. It is argued here that fusion processes would be better equipped to produce these entities estimates if designed in conjunction with formal ontologies. Such ontologies can provide fusion system designers with a comprehensive, and computationally tractable, phenomenological description of a given domain. A properly conceived formal ontology is necessary to provide a structure for analyses of domain-specific objects, object attributes, and relations and to assure both interoperability and reusability of the designed fusion system for different domains. In this sense, a formal ontology *informs* the fusion process by providing a phenomenological description of a given state of affairs. At the same time the fusion process *constrains* the formal ontology in terms of its size, scope, and relevance for certain decision-making needs. This report presents ontology related concepts for building a threat ontology (ThrO) to be used in various information fusion applications involving threats.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Information fusion, formal ontology, threat.

Defence R&D Canada

Canada's Leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale

DEFENCE DÉFENSE

www.drdc-rddc.gc.ca