

The Semantic Battlespace Infosphere: A Knowledge Infrastructure for Improved Coalition Inter-Operability

Paul R. Smart and Nigel R. Shadbolt

School of Electronics and Computing Science, University of Southampton, Southampton, SO17 1BJ, UK
{ps02v, nrs}@ecs.soton.ac.uk

Abstract—Operational effectiveness in coalition environments is based on the need for inter-operability at a variety of levels. While inter-operability concerns are most easily thought of in terms of technology, this paper emphasizes the importance of consensual interpretations of the semantic significance of exchanged information. In this paper we outline some of the challenges to effective modes of information exchange in coalition operational contexts. We also discuss potential approaches to these challenges in the context of a semantically-enabled technological framework for information exploitation – the Semantic Battlespace Infosphere (SBI). Relevant aspects of this framework are introduced and some of the socio-technical challenges that are likely to be encountered are discussed.

1. INTRODUCTION

Both the UK and the US recognize the importance of transnational alliances as the basis for future military operations. The vision is that military activities spread across all levels of the operational spectrum (from large-scale war-fighting to peace support and humanitarian assistance) will assume the form of ‘coalitions of the willing’ [1]. Such coalitions will demand the close inter-operation, but not necessarily integration, of multi-national forces, each of which brings its own set of technological, ideological, organizational, procedural and cultural idiosyncrasies to the theatre of operations. Issues of inter-operability therefore sit at the heart of many research efforts concerned with the future effectiveness of military coalitions. Inter-operability is particularly important in contexts that involve non-military agencies, e.g. diplomatic, humanitarian and civil authorities. In addition, the rapid, opportunistic exploitation of situation contingencies, the need to self-synchronize and the requirement to synergistically marshal diverse military assets in the context of agile force structures, requires the ability to exploit and share information in ways that transcend the traditional boundaries of national affiliation and operational environment (i.e. land, sea and air).

This paper discusses some of the issues related to improved coalition inter-operability that are being researched in the context of the International Technology Alliance¹ (ITA) program. Our approach to coalition inter-operability is grounded in the use of Semantic Web technologies and is organized around a framework for advanced modes of

information integration, exploitation and exchange in coalition military contexts. We refer to this framework as the Semantic Battlespace Infosphere (SBI) and present it as an extension to an existing information management system known as the Joint Battlespace Infosphere (JBI) [2, 3]. This paper summarizes the basis for the SBI and describes some of the socio-technical issues that may limit the acceptability or viability of proposed solutions.

2. COALITION INTER-OPERABILITY

Inter-operability issues are most easily thought of in terms of technology, but they can actually assume a variety of forms. Inter-operability issues may, in fact, arise on a number of levels, including, the technological, the socio-cultural, the operational and the epistemic or knowledge-based. While this paper does not attempt to belittle the importance of any of these levels, we argue that many of the most difficult problems, with regard to inter-operability, lie in the realm of ‘epistemic inter-operability’, i.e. the compatibility that exists between coalition partners with respect to conceptual models, inference processes, reasoning strategies, etc. In these cases, our concern is not merely limited to issues regarding the physical exchange of information; we also have to consider the meaning assigned to information items and the interpretive biases possessed by coalition partners. Such factors can influence the interpretation of information content in terms of both its semantic referents (what the information is about) and its semantic significance (the implications of the information for current and future action). Issues of epistemic inter-operability also extend beyond simple information exchange contexts; they subsume the ability to integrate and fuse information from physically disparate and semantically heterogeneous information sources, as well as the ability to coordinate and orchestrate the activities of distributed problem-solving agents. A key element of these abilities, we argue, lies in the capacity to exploit semantically-enriched representations and to establish mappings between ostensibly disparate representations by virtue of their semantic similarity. The framework we propose, the SBI, presents one strategy for the realization of these semantically-enabled capabilities.

3. SEMANTIC BATTLESPACE INFOSPHERE

The JBI is a combat information management system designed to acquire and represent information from a wide

¹ <http://www.usukita.org/>

variety of information sources with the express purpose of supporting enhanced situation awareness, coalition inter-operability and operational flexibility [4]. The JBI exploits information available in existing C2 (Command & Control) systems, but it does not aim to replace them. Rather, the emphasis is on information exchange between these systems. The JBI also aims to support improved situation awareness based on its ability to integrate (fuse) information from different sources and to make inferences based on environmental data. Situation awareness is further supported by the proposed ability of the JBI to tailor information content to suit the needs of individual end-users: “the commander gets high-level coverage of the campaign, while the soldier in the field gets a detailed description of a nearby hostile base” [2; pg. iii].

It is not the purpose of this paper to review the relative merits or demerits of the approach advocated by the JBI. We suggest that the basic vision proposed by the JBI is sound, but that its applicability to coalition inter-operability is undermined by its failure to fully embrace semantically-enriched representational schemes. Such schemes would seem indispensable for a number of reasons, not least because of the difficulty of agreeing a common language or vocabulary that is accepted across all elements of a coalition formation. Even if a common vocabulary could be agreed, it would not necessarily extend to the variety of humanitarian, diplomatic and government agencies with which coalition partners must often inter-operate. Moreover, semantically-enriched forms of representation support the easy revision and maintenance of existing vocabularies and conceptual models. New terms, such as ‘Qandahar’, can easily be added and associated with existing terms, such as ‘Kandahar’, without causing major disruption to existing capabilities. To address the potential shortcomings of the JBI we propose the notion of the SBI, which attempts to take the basic vision of the JBI and extend it with respect to semantically-enabled capabilities².

The notion of the SBI does not refer to a specific system or capability so much as a framework for inter-operability-focused modes of technology exploitation and development. Like the JBI, the SBI attempts to provide a capability for inter-operability between previously disparate C2 systems; it also aims to facilitate situation awareness with regard to information integration and publish/subscribe mechanisms. Unlike the JBI, however, the technological commitments endorsed by the SBI concept differ from those characterizing the JBI initiative. In particular, the notion of the SBI countenances a framework within which a variety of semantic technologies are used to support key capabilities (e.g. meaning-preserving modes of information exchange, enhanced situation awareness, a common understanding of the significance of events and information items, etc.) by virtue of their capacity to support and exploit semantically-enriched representations within highly distributed network environments. It builds on many of the guiding principles

and technologies of the Semantic Web initiative [5], but the aim is to adapt and extend these principles and technologies to suit the military environment.

The core technological ingredients of the SBI are domain ontologies, created using languages such as the Web Ontology Language (OWL) [6, 7] and Resource Description Framework (RDF) [8]. These ontologies constitute the representational bedrock for inter-operability-related capabilities such as mediation services between federated systems of systems, meaning-preserving modes of information exchange, and query execution at semantic levels of abstraction. Although the technological approach we advocate relies on the use of Semantic Web languages such as RDF, RDFS and OWL, the choice between these languages is often a trade-off between the relative simplicity of languages such as RDF (which supports ease of development and maintenance) versus the greater semantic expressivity of languages such as OWL (which supports greater reasoning capabilities). Although simplicity is to be countenanced in situations where we want to accomplish the widespread adoption of a representational language, it is questionable whether RDF will support the kinds of capabilities we desire in respect of coalition inter-operability solutions. Such solutions are founded on the ability to identify semantic correspondences between ostensibly disparate vocabularies [9] (a capability variously referred to as semantic integration, ontology mapping or ontology alignment), but languages such as RDF may not provide enough semantic constraints in order to effect this mapping (at least not one that can support wholly automated alignment solutions). Ideally, what we seek to accomplish with respect to ontology mapping, is an automated approach that capitalizes on the capacity of semantically-enriched representational formalisms to denote the semantic similarity between ontology elements (i.e. classes, properties and instances). With its greater degree of semantic expressivity, OWL is much better equipped to provide this capability, but a reliance on OWL may affect the tractability of a knowledge capture solution (see Section 4.1), especially in a coalition environment where the extent of the conceptual space and the cultural heterogeneity of participating agencies may undermine any effort to develop a single, all-encompassing ontology.

Semantic queries are another essential element of the SBI framework. Semantic query languages, such as SPARQL, work in conjunction with ontologies to increase the accessibility of information content in a way that is more likely to address the goal-relevant epistemic needs and requirements of a particular problem-solving agent. The point is that unlike conventional query languages, such as SQL, semantic query languages operate at the semantic level – they focus on the conceptual structure of the domain and they do not make any assumptions about the underlying structure of the dataset used to store data within the domain. This shift in focus is of potential benefit because it allows queries to be expressed in a form that makes sensible contact with the kinds of conceptual abstractions and

² Note that our aim here is not to argue for the replacement of the JBI, so much as its extension with semantically-enabled capabilities.

relationships identified within a domain of discourse³.

Languages such as SPARQL provide an effective semantic query solution, and many knowledge repositories cater for such queries, e.g. 3Store [10]; nevertheless, it is not entirely clear how such queries should be executed in an environment where knowledge content is distributed across multiple nodes of a dynamic (and sometimes *ad hoc*) network environment. One particular problem concerns the likely volatility of semantic query results in military contexts characterized by *ad hoc*, mobile and wireless communication infrastructures. The key problem is that time-variant changes in network connectivity (or the differences in connectivity apparent from the perspective of physically distributed military agencies⁴) results in the differential availability of nodes and their associated knowledge resources. This can contribute to a confusing situation picture because query results executed from one location in the network need not coincide with the results of the same query executed elsewhere. Moreover, the same query may return different results at different times based on the physical distribution of knowledge resources and the extent of intervening changes in network topology. The distributed nature of knowledge resources is a potential problem here because it complicates the possibility of establishing a common collective representation about the nature and implications of the current situation picture. Ultimately, we argue, this can attenuate shared situation understanding and situation awareness and undermine the potential for coalition inter-operability.

These concerns about query execution in distributed environments also apply, to some extent, to reasoning processes. To the extent that reasoning processes subtend multiple, physically distributed resources, then reasoning outcomes will depend on the relative stability of the network infrastructure across multiple invocations of the same reasoning process (either from different points in the network or from the same point at different times). As with query capabilities, the dynamic nature of the military network environment (in contrast, perhaps, to the situation with the World Wide Web), potentially undermines the possibility for a common understanding of the operational picture, especially when coalition elements are geographically distributed and connectivity privileges are non-uniform.

³ One concern, of course, from an inter-operability perspective is that such languages, despite exploiting domain ontologies, are likely to be largely ineffectual in a situation where different coalition partners are making distinct ontological commitments (i.e. using different, albeit overlapping, ontologies within the same problem domain). However, such concerns can be allayed, to some extent, by the use of ontology alignment solutions. In this case, semantically-equivalent entities in separate ontologies can be targeted by the query, despite the fact that they may be identified using different linguistic labels.

⁴ For example, imagine a sub-component of the larger network infrastructure becomes temporarily isolated or has limited connectivity to the remainder of the network. In this case, the networked resources available to the same query executed at different points in the network could yield radically different query results.

Like the JBI, the SBI does not aim to replace existing systems; rather the idea is that the SBI will act as the middleman within a federated system of systems. The SBI thus aims to serve as a mediator between previously disparate systems, enabling applications and services to exchange information in ways that preserves the original semantics of information content. Note that this does not mean that information will necessarily remain the same (e.g. with respect to its physical form) across information exchange contexts. The point about using domain ontologies, in conjunction with ontology alignment solutions, is that the emphasis is on the semantics of the information content, not the information content *per se*. Thus, if information was required to be transformed in the course of an exchange involving culturally-disparate user communities, then such a transformation would be undertaken in an effort to avoid semantic ambiguity and misinterpretation by the target community. The key point here is that our mediation solution does not simply aim to provide a mechanism for information exchange; rather the emphasis is on enabling information exchange with respect to common semantic frames of reference, frames of reference that make explicit the meaning of information content to all coalition partners.

The emphasis on explicit semantics and the capability for flexible modes of information transformation to support meaning-preserving modes of information exchange is something which tends to be overlooked by many other approaches to information exchange in the military domain. Conventional IEDM solutions [11], for example, may be poorly suited to meaning-preserving modes of information exchange simply because such approaches do not avail themselves of a sufficiently rich repertoire of semantically-enriched representational formalisms. The JBI is also deficient in this respect. The definition of battlespace objects in the JBI is based on XML Schema language, but such languages are potentially inadequate in a coalition environment where data standards may be difficult to establish, enforce and sustain. Perhaps most importantly, an XML schema provides a syntactic specification that defines the structural organization of data. While this is valuable as a mechanism for data exchange, it does not provide any means for the effective representation of semantic information, i.e. what the data means in terms of its actual relationships to other data within the same dataset and potential relationships to data that may be defined elsewhere and received at a different time. To capture the semantic significance of data requires knowledge about classes of data objects and how these objects relate to one another. This type of information is precisely what domain ontologies within the SBI framework are intended to capture.

4. (SOCIO-)TECHNICAL CHALLENGES

The realization of capabilities implied by the notion of the SBI depends on the resolution on a number of key challenges. Some of these challenges derive from the idiosyncratic nature of the military network environment,

especially the prevalence of *ad hoc*, mobile, wireless, networks; others relate to more general issues that pervade the Semantic Web and user-centred design communities. Above all we suspect that many of these challenges encompass the social, cultural and psychological domains - they are not merely related to the provision of a technological and representational substrate for meaning-preserving modes of information exchange. The point is that we may need to pay careful attention to the cultural differences between coalition partners and the organizational, psychological and social contexts that influence the acceptability and/or usability of proposed solutions. The following sections provide an overview of (some of) the potential challenges we face with respect to a realization of semantically-enabled capabilities supported by the SBI.

4.1. Knowledge Capture

A key limiting factor in the widespread adoption and use of ontologies is the overhead associated with their initial development and subsequent maintenance. Ontology development relies on the capture and formalization of domain knowledge, and this can sometimes serve as a significant bottleneck in the knowledge engineering process [12]. Even when individuals or institutions can marshal the resources to overcome this bottleneck, problems of completeness (does the ontology achieve adequate coverage of the target domain?), consistency (does the ontology align itself with the representational strategies adopted by other agencies?) and topicality/currency (does the ontology reflect the current conceptual focus and complexity of target user communities?) all threaten to undermine the potential utility of an ontology.

One approach to overcoming these problems focuses on the extent to which some aspects of the semantic infrastructure of a domain can be derived from the actions and interactions of agent communities within the domain. The idea here is that information networks serve as a substrate for patterns of activity that contribute to the (automatic) pooling of knowledge and expertise within specific communities of interest. One example, of this phenomenon is represented by the notion of ‘collaborative filtering’ wherein, for example, patterns of consumer activity associated with the purchase of books, CDs, DVDs, etc., allows for the generation of product recommendations that then influence subsequent purchase patterns [13]. Another example is provided by customizable tagging schemes, such as those used by Flickr. Such schemes devolve much of the responsibility for initial ontology (or at least taxonomy) development to the user community, which avoids the upfront cost for agreeing upon a taxonomy when, perhaps, the nature of the information to be collected and its use are not yet known. It also allows the taxonomy to emerge and change dynamically as additional information is accumulated. One aspect of our research in the ITA aims to explore these and other techniques as a means for understanding the extent to which networked patterns of activity can be exploited to infer or derive some insight into the semantic infrastructure of a domain.

4.2. Distributed Knowledge Infrastructure

A second major challenge to semantically-enabled capabilities, in the context of the SBI, is the notion of a distributed knowledge infrastructure. Just as the vision of the Semantic Web is a vision of an extended Web of machine-readable information and automated services [5], so the notion of the SBI builds on the capabilities (and limitations) of existing military information and communication infrastructures – infrastructures in which knowledge (and the services they support) will often be distributed across multiple nodes of a large-scale information network. Decentralization of resources within such networks contributes to the resilience of higher-order capabilities in the face of network disruption (i.e. the failure of network elements), but it also implies a number of distinct technical challenges, including (but certainly not limited to) the following⁵:

1. How are we to orchestrate and coordinate the activity of services in a dynamic network environment, an environment in which different service-related capabilities are dependent on the time-variant topological organization of the network?
2. How are we to fully exploit knowledge and services in an environment with no centralized standards agency? For example, how are we to deal with the tendency of novel coalition elements to describe knowledge and services in a manner that reflects their idiosyncratic (perhaps culturally-entrenched) modes of conceptualizing problem domains and describing service-capabilities?
3. How can we cope with the various performance constraints associated with network infrastructures (e.g. the latency required to aggregate knowledge across multiple network nodes) to yield knowledge services (e.g. reasoning capabilities) that complete within an operationally useful timeframe?

There are a variety of potential responses to these challenges, not all of which fall within the scope of the current paper. Tentative solution strategies include the active (knowledge-driven?) configuration of network topologies to better suit specific knowledge processing requirements, the duplication of knowledge resources throughout the network in order to support a form of epistemic redundancy in the face of network dynamics, and the use of query caching strategies to achieve virtual knowledge-processing stability in the face of network disruption. However, many of the inter-operability challenges to be tackled in the context of our SBI research programme relate to the need to align disparate ontologies and integrate information content from multiple semantically-heterogeneous information sources. This is the challenge of semantic integration.

⁵ These are in addition to the problems associated with the potential volatility of semantic query results as discussed in Section 3.




Figure 1: Ontology mapping/alignment

4.3. Semantic Integration

As the global information environment becomes increasingly pervasive and spans ideologically, culturally and ethno-linguistically diverse communities, so the information exchange challenge becomes ever harder. The battlefield network currently consists of many information sources, including in-theatre sensors, platforms, intelligence reports and remote information such as archival intelligence and satellite data. In future coalition contexts, strategies for operationally-effective modes of information exchange and exploitation will need to target a wider variety of disparate information repositories and communication systems, including digital datalinks, military information repositories, and the totality of the information space available via internet-enabled and peer-to-peer computing environments. In situations such as these the potential for semantic ambiguity is rife because the meaning of symbolic information often reflects the experiential, epistemic, cultural and task-specific biases of the information provider. Both the semantic referents and semantic significance of information is not invariant with respect to information exchange contexts, rather one sees a degree of semantic specificity - a community specific interpretation of meaning that may not necessarily transcend cultural, organizational and/or national boundaries. The point is that once we encounter distributed network environments that subtend a wide variety of information systems, sources and user communities (as is often the case in coalition contexts), we face a critical challenge in terms our ability to integrate and share information in a semantically-sensible manner (one that respects the meaning assigned to information content by the originating agent or agency). We refer to this as the semantic integration challenge.

Our approach to this challenge, in the context of the ITA, is grounded in the use of semantically-enriched domain knowledge models (ontologies) and the notion of ontology alignment. Ontology alignment (see Figure 1) is a key element of semantic integration [9]. Its aim, in essence, is to establish mappings between the elements of multiple (ostensibly disparate) domain ontologies as a means of identifying semantically-equivalent sub-components. Once established, these mappings can be used to drive information aggregation and integration activities (as might be required for goal-relevant information processing [see 14], for example), but most of all we see such mappings as a mechanism for ontology-based mediation of information

exchange between conceptually or linguistically disparate communities.

A wide variety of tools and techniques have emerged to support ontology alignment [9, 15] all of which rely, to a greater or lesser extent, on the similarity between ontology elements, e.g. in terms of common relationships, linguistic labels or instance sets [15]. Nevertheless, there is fertile ground for research here both in terms of the accuracy of the techniques and the degree of automation that each technique supports. One potentially interesting strategy is to explore cognitive science techniques for concept mapping, which, in general, have received little attention in the Semantic Web community. Goldstone et al [16] thus describe a technique that is grounded in the use of neural networks to establish mappings between the elements of a conceptual system. Such approaches, they argue, could be successfully applied to problems in ontology alignment.

One aspect of our approach in developing the SBI, and one that most strongly discriminates our work from previous work in ontology mapping/alignment, concerns the need to dynamically integrate and align ontology fragments in specific task contexts (see Figure 2). In essence this approach eschews the idea of large-scale, global ontology alignment independent of task context; rather it countenances the idea that ontologies (or relevant sub-components thereof) should be dynamically aligned to reflect goal and task-relevant processing. Such capabilities rely on effective mechanisms to represent the epistemic requirements of tasks, and to prune larger ontologies in light

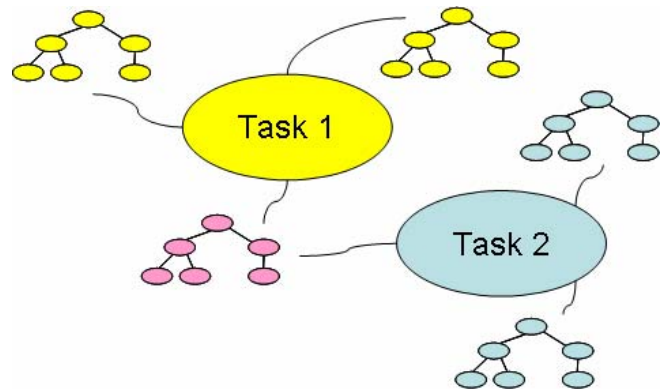


Figure 2: Task-specific semantic integration

of these requirements, i.e. to extract just those elements that are relevant for current problem-solving activity.

4.4. Visualization & Interaction

The human end-user is a key element of the SBI. Even in the case of automated services, the outcome of such services needs to be carefully geared to suite the needs and requirements of the human operator in relation to problem-solving objectives. The SBI is, in essence, intended to provide a framework for assisting coalition members with respect to knowledge processing, and thus its success is ultimately predicated on the ability of application interfaces to support operationally-effective modes of working (and

thinking).

User-centred design approaches are likely to be of paramount importance in understanding the opportunities for operationally-effective information exploitation within coalition contexts. One approach we have adopted in the ITA is the notion of Goal-Directed Task Analysis [14]. Such an approach embraces both user-centred design principles and also focuses on the goals of the human operator. In essence it provides a framework for understanding the information requirements of the decision-maker in specific operational contexts and helps to drive information aggregation, semantic integration and service coordination processes in respect of goal-oriented information processing. Such analyses are critical in terms of understanding the (sometimes subtle) differences in problem-solving strategies adopted by different coalition partners, and they help to customize semantically-enabled information aggregation and integration processes in ways that best support shared situation awareness and collaborative problem solving [14].

5. CONCLUSION

In this paper we have attempted to highlight the potential role played by semantic technologies with regard to coalition inter-operability issues. We have introduced the notion of the SBI, which provides a vision of semantically-enabled capabilities in the future battlespace environment, and suggested ways in which this framework could facilitate inter-operability in military coalition contexts. The aim of the SBI is, in essence, to develop a common semantic frame of reference to support consensual interpretations of entities, events and actions across force elements and between coalition partners. This is important because future coalition capabilities will critically depend on an ability to exchange information in ways that preserve the meaning assigned to information content, especially when the operational context demands close cooperation with non-military agencies (e.g. diplomatic, humanitarian and civil authorities in the case of peace support and humanitarian relief contexts). A key objective of the SBI is therefore to provide a foundation for coalition inter-operability, enabling semantic integration with respect to both digital datalink networks and unstructured, non-military information sources. Notwithstanding the ready availability of extant semantic technologies (e.g. OWL, RDF, SPARQL, etc.), a number of aspects of the military coalition environment make the solution space somewhat different from that which we encounter in the case of the Semantic Web. Such challenges include the need to deal with the idiosyncrasies of the military network infrastructure (e.g. its mobile, *ad hoc* and wireless nature), the need to engender effective information exchange solutions via ontology alignment/mapping mechanisms and the need to deal with the potential (perhaps culturally-entrenched) differences between coalition partners with respect to the interpretation of situation-relevant information. These issues impact directly on coalition inter-operability issues and form a key component of our ongoing research in the ITA.

ACKNOWLEDGMENT

This research was sponsored by the US Army Research laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. The authors would like to express their thanks to Dave Braines at IBM UK Ltd for his comments on an earlier draft of this paper.

REFERENCES

- [1] "Delivering Security in a Changing World: Defence White Paper," Ministry of Defence, UK 2003.
- [2] "Building the Joint Battlespace Infosphere: Volume 1 - Summary," United States Air Force - Scientific Advisory Board, USA SAB-TR-99-02, 2000.
- [3] "Building the Joint Battlespace Infosphere: Volume 2 - Interactive Information Technologies," United States Air Force - Scientific Advisory Board, USA SAB-TR-99-02, 1999.
- [4] "Information Management to Support the Warrior," United States Air Force - Scientific Advisory Board, USA SAB-TR-98-02, 1998.
- [5] T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific American*, vol. 284, pp. 34-43, 2001.
- [6] G. Antoniou and F. van Harmelen, "Web Ontology Language: OWL," in *Handbook on Ontologies*, S. Steffen and R. Studer, Eds. Berlin: Springer Verlag, 2003.
- [7] D. L. McGuinness and F. van Harmelen, "OWL Web Ontology Language Overview," World Wide Web Consortium, 2004.
- [8] B. McBride, "The Resource Description Framework (RDF) and its Vocabulary Description Language RDFS," in *Handbook on Ontologies*, S. Steffen and R. Studer, Eds. Berlin: Springer Verlag, 2003.
- [9] Y. Kalfoglou and M. Schorlemmer, "Ontology mapping: the state of the art," *Knowledge Engineering Review*, vol. 18, pp. 1-31, 2003.
- [10] S. Harris and N. Gibbins, "3store: Efficient Bulk RDF Storage," presented at 1st International Workshop on Practical and Scalable Semantic Web Systems, Sanibel Island, Florida, USA, 2003.
- [11] E. Lasschuyt, M. van Hekken, W. Treurniet, and M. Visser, "How to make an effective information exchange data model," NATO RTO IST-042 Symposium on Coalition C4ISR Architectures and Information Exchange Capabilities, The Hague, Netherlands JWP 3-52, 2004.
- [12] F. Hayes-Roth, D. A. Waterman, and D. B. Lenat, *Building Expert Systems*. Reading, Massachusetts: Addison-Wesley, 1983.
- [13] G. Linden, B. Smith, and J. York, "Amazon.com Recommendations: Item-to-Item Collaborative Filtering," *IEEE Internet Computing*, vol. 7, pp. 76-80, 2003.
- [14] M. Endsley, B. Bolte, and D. G. Jones, *Designing for Situation Awareness: An Approach to User-Centered Design*. London: Taylor & Francis, 2003.
- [15] "State of the art on ontology alignment," University of Karlsruhe, Karlsruhe, Germany Knowledge Web Deliverable 2.2.3, 2004.
- [16] R. L. Goldstone, Y. Feng, and B. J. Rogosky, "Connecting concepts to each and the world," in *Grounding Cognition: The Role of Perception and Action in Memory, Language and Thinking*, D. Pecher and R. Zwaan, Eds. Cambridge, UK.: Cambridge University Press, 2005.