

Knowledge-Based Information Fusion for Improved Situational Awareness

Paul R. Smart
ps02v@ecs.soton.ac.uk

Nigel R. Shadbolt
nrs@ecs.soton.ac.uk

Leslie A. Carr
lac@ecs.soton.ac.uk

monica c. schraefel
mc@ecs.soton.ac.uk

Department of Electronics & Computer Science
University of Southampton
Southampton, SO17 1BJ, UK

Abstract - *This project aims to investigate knowledge-based approaches to the problem of information fusion. We show how domain-specific knowledge, in conjunction with semantic web services, can be used to improve situational awareness in the operational deployment of humanitarian operations, especially when such operations occur against a contextual backdrop of ongoing military conflict. Information, harvested from a variety of physically disparate and semantically heterogeneous data sources, is used to establish a common repository of semantically-rich, conceptual-level representations of real-world events. Reasoning services, deployed over the knowledge repository, are used to intelligently fuse information, taking into account factors such as the level of trust and confidence assigned to specific information sources. Selective attention to information of relevance to particular users helps avoid situations of information overload that might otherwise result from the information fusion process. We believe this approach to have general applicability in a wide variety of military contexts.*

Keywords: decision support, humanitarian relief, information fusion, knowledge-based systems, ontologies, semantic web, situational awareness

1 Introduction

Increasingly, the activity of UK military forces occurs against a backdrop of humanitarian intervention. In some cases the need for humanitarian intervention relates to the direct affects of military activity itself, while in other cases military activity may be sanctioned on humanitarian grounds. In either case, cooperation with humanitarian agencies and improved awareness of the temporal unfolding of humanitarian events is increasingly critical to the success of modern military forces. Even when major hostilities have ceased, or fail to develop, there often remains a requirement for military support to bolster the efforts of humanitarian aid workers, either with respect to logistics, security or aid distribution.

A central concern in the coordination of both military and humanitarian operations relates to enhanced situational awareness. With the widespread adoption of the Internet as a communication medium, the continued advance of the mass media and the increasing availability of satellite and cellular communications technology, the lack of information is seldom a problem in today's information age. Rather the problem is one of information overload and its solution lies in the ability to filter information in a manner that befits the knowledge and information processing objectives of key knowledge workers. Fortunately, a range of technologies are emerging as part of the Semantic Web initiative [1] that may help to attenuate the problem of information overload by providing more expressive mediums for the representation of information content. In the current project we aim to show how Semantic Web technologies can be used to improve situational awareness and operational effectiveness in the planning, coordination and deployment of humanitarian aid operations, especially when such operations occur against a backdrop of hostile military activity. Central to improved situational awareness in these contexts, we argue, is the ability to fuse information from a variety of physically disparate and semantically heterogeneous sources in order to establish a common picture of real-world events and actions. Semantic Web technologies, we argue, are invaluable with respect to this fusion-related capability.

2 System Functionality

To showcase the role played by knowledge-related technologies in improved situational awareness via intelligent information fusion, we have developed a Technical Demonstrator System (TDS), which integrates a variety of Semantic Web Services (SWSSs) within a common application framework. The primary objective of the TDS is to highlight how real-time information, harvested from a variety of physically disparate and semantically heterogeneous information sources, can be interpreted with respect to formal ontological characterizations of the problem domain in order to improve operational effectiveness and situational

awareness in the domain of humanitarian aid. Subsequent sections provide a functional overview of the system with respect to a number of capability areas.

2.1 Information Retrieval

Information retrieval subsumes the ability to receive, monitor and actively acquire information from a variety of information sources, e.g. web pages, online databases, tactical datalinks, email notifications, etc. In some cases the TDS is passive with respect to the receipt of this information (for instance it is notified of the occurrence of certain events, e.g. an earthquake); in other cases information is actively harvested from information sources. Information harvesting, in this context, entails the extraction of pertinent information from general information sources, frequently adopting natural language techniques to mine documents or web pages. Harvesting can occur on a scheduled basis, e.g. once a week, or in response to the epistemic requirements of the system. For example, the system may seek supporting information regarding the occurrence of a particular event in order to increase its confidence that the event has actually occurred.

2.2 Information Triage

Information triage concerns the ability of the system to assess the semantic relevance of received information and deliver filtered report streams to different user groups in a timely manner. Central to this ability is the notion of agent characterizations that indicate the information and knowledge requirements of different user groups. Agent characteristics such as task context, organizational affiliation and executive role play an essential role in the system's ability to evaluate the relevance of different information items and to focus attention on those information items of direct relevance to an agent's ongoing needs and concerns.

Information triage is also undertaken with respect to the information directly received by the system. In this case the emphasis is on assessing the semantic relevance of incoming information streams with respect to the goals and objectives of the system itself. As with the dissemination of information to external agents, unwanted information is filtered from the information stream; however, in the former case the information may still be processed by the system, whereas in the latter case the information is considered irrelevant to the system's activity and is discarded. The goal in both cases is to focus attention on a subset of the total information available to an agent thereby optimizing the distribution of cognitive and perceptual resources with respect to information load.

2.3 Information Fusion

Information fusion concerns the ability of the TDS to integrate and make sense of a variety of information inputs in a manner that takes into account the relative reliability, accuracy and provenance of information in order to build a coherent representation of events and objects in the immediate operational environment. Central to this fusion-related capability is the notion of trust. The level of trust of invested in each of the incoming information streams determines the degree of confidence the system has in the occurrence of particular events and their associated characteristics. In extreme cases, where information is received from an entirely untrustworthy source, the system may simply fail to 'believe' any of the information received from the source. Typically, however, the system will assign varying degrees of confidence to the information received from different sources based on the level of trust invested in each source.

Every information item encountered by the system is associated with meta-data that expresses the level of certainty or confidence adopted by the system with respect to the information item. This includes information derived from the inferences made by the system itself, in which case the certainty associated with the inferred information is based on the certainty assigned to each of the information items on which the inference was originally made. Certainty information is thus propagated along complex chains of inference execution ultimately impacting on the manner in which different decision-outcomes are sanctioned by the TDS.

2.4 Knowledge Processing

Once a coherent representation of the operational environment has been established, the system may engage in a number of additional knowledge-related processes. These processes are primarily geared towards the progressive refinement and enrichment of the conceptual model established by fusion-related activities, especially to the extent that such refinements yield information of direct relevance to the coordination, planning and deployment of humanitarian relief efforts. An example concerns the ability of the system to infer the scale of a humanitarian disaster when provided with information about the occurrence of a geological event in the outside world, e.g. an earthquake. Based on knowledge about the magnitude of the event (e.g. Richter scale value), the chain of events in which the event is either the cause or a consequence (e.g. an earthquake of a given magnitude may cause structural collapse or tsunami events depending on its epicentral location), and the population density of regions affected by the event (or its causal precursors or sequelae), the system should be capable of predicting the location, scale and type of impending humanitarian disaster. Once predicted, the system should

be capable of determining the type of humanitarian intervention required, e.g. the type of aid resources required by the disaster, based on knowledge about the resource requirements of different types of events. Finally, it should be capable of automatically notifying humanitarian (and other) agencies that have registered an interest in the occurrence of this particular type of event based on its typology, location, resource requirements and other criteria.

2.5 Information Dissemination

The ability of the system to enhance situational awareness and operational effectiveness in the deployment of humanitarian relief efforts requires that information and knowledge ultimately be communicated to executive agencies, i.e. those capable of implementing real-world actions. The key function of the system in this respect is to augment, rather than replace, the cognitive and deliberative faculties of the agents with whom it inter-operates, typically by notifying and alerting executive agencies of the occurrence of particular events or information items upon which their own particular decision-making processes are based. In essence we see this role of the system as conforming to the traditional vision of a knowledge management system whose principal aims are to disseminate (selected) information to the right agencies in a timely fashion in a format suited to their idiosyncratic perceptual and cognitive profiles. The notion is, of course, encapsulated in the knowledge management mantra of systems aimed at getting *the right knowledge to the right people in the right form at the right place at the right time*. In order to fulfill this information dissemination function the system requires background knowledge about the roles of external agents, their executive capabilities, responsibilities to act, position in power hierarchies and patterns of communication with other agents. The latter of these agent characteristics is perhaps of crucial importance since the tasks undertaken by external agents will themselves typically occur in a rather distributed fashion in which the different capabilities of particular agents are recruited to implement distinct parts of a larger task. In essence we regard it as necessary for the system to have some notion of the ‘community of practice’ of each registered agent.

2.6 Interaction and Visualization

The way in which information is presented to agents can exert a profound impact on their subsequent information processing activities and decision outcomes. Such considerations are relevant to the current initiative inasmuch as we wish to maximize situational awareness and optimize information transfer to executive agents without producing a situation of information overload. In some cases the types of visualization and modes of

interaction will be invariant with respect to specific user groups. Such will be the case when military authorities or codes of practice sanction the use of particular types of user interface for particular types of user group, or when the exploitation of interface structure has become in an intrinsic and inextricable part of expert performance. In other cases more flexible strategies of information transfer can be adopted and these need to take into account the specific cognitive and perceptual preferences and biases of end users. While, our investigation into optimal user interface design is yet to be completed in the context of the current initiative, we aim to countenance the selection of those information-bearing environmental structures that best support operationally effective modes of working and thinking.

In addition to the cognitive and perceptual biases of end user agents we also need to consider how the display capabilities of different communication devices impacts on interface design issues. Clearly, the types of visualization supported by a high-resolution, flat panel, color display at command headquarters will not easily transfer to a handheld PDA device used by a field operative.

3 Ontology Development

The key to success in each of the capability areas detailed in Section 2 is the provision of formal ontological characterizations of the target problem domain. Recently, attempts to provide a set of representational formalisms for the communication of ontological structures within the framework of the Semantic Web have coalesced around the Web Ontology Language (OWL) initiative [1-3]. We chose OWL as the preferred medium for representing knowledge in the context of the current initiative for the following reasons:

1. RDF compliance: OWL can be serialized in an RDF-compliant format and is therefore compatible with a range of extant semantic web technologies built around RDF.
2. Tool support: A number of knowledge editing environments provide direct or indirect support for OWL, e.g. Protégé [4, 5]. In addition, some widely available APIs, such as Jena [6], provide support for querying OWL-based repositories.
3. User Community: OWL has gained widespread acceptance amongst the academic community and is currently a key component of a number of major Semantic Web research initiatives, e.g. myGRID [7, 8], AKT [9, 10].
4. Semantic expressivity: OWL supports a high degree of semantic expressivity by virtue of its background in description logics. Such expressivity easily lends itself to certain forms of automated reasoning, particularly the

subsumption reasoning undertaken by DIG-compliant (DL Implementation Group) reasoners such as RACER [11].

Our approach to ontology development in the context of the current initiative assumed the form of a tripartite scheme in which the knowledge outputs of an earlier phase were progressively refined and formalized by subsequent phases. Each of the phases of our approach is detailed in subsequent sections.

3.1 Knowledge Capture

The first phase relied on the careful use of a number of knowledge acquisition techniques [12, 13], e.g. protocol analysis, ladder grids, repertory grids, etc., to acquire and elicit knowledge from a range of source materials, mostly public domain resources such as websites. PCPACK, a tool suite for software-assisted knowledge acquisition (marketed by Epistemics [14]), was used extensively in this early phase of the knowledge engineering life cycle. The outcome of this particular phase of the knowledge engineering initiative consisted of largely semi-structured characterizations of domain knowledge: annotated textual sources, simple taxonomic hierarchies, process models and so forth.

3.2 Knowledge Modeling

This phase of the knowledge engineering process entailed the development of detailed knowledge-level models of the domain conceptualizations, knowledge-rich contingencies, inferential processes, and so forth, that are required for successful problem-solving in the target problem domain. For this phase of the knowledge engineering life-cycle we relied on the knowledge model specification provided by the CommonKADS methodology [15]. In practice, we have found that the level of representational detail afforded by the CommonKADS Knowledge Model is ideal in terms of enabling communication between stakeholder groups without reneging on the commitment to provide detailed, machine-readable representations of domain knowledge.

3.3 Ontology Formalization

The final phase of the knowledge engineering life-cycle involved the transformation of CommonKADS Knowledge Models to OWL-based ontologies. To some extent, the transformation was accomplished automatically, although in some cases the use of description logic expressions by OWL prohibits automatic translation from largely frame-based models. Our strategy, in this case, was to perform a partial translation of the knowledge structures and then edit the partially completed models with an OWL-based editor, namely Protégé with an OWL plug-in [4].

Interestingly, some features of the CommonKADS Knowledge Model could not be easily accommodated in

the OWL model, irrespective of manual intervention. These features relate to the representation of knowledge-rich contingencies in the form of IF-THEN or production rules. The CommonKADS Knowledge Model adopts a representational strategy whereby these contingencies are represented as expressions in the antecedent and consequent of an implication rule structure, but OWL does not, at the present time, incorporate an analogous representational device. While efforts are underway to address this issue in the form of the Semantic Web Rule Language (SWRL) [16], a combination of OWL and a dedicated rule implementation language called RuleML [17], we opted, in the context of the current initiative, to maintain a separate representation of rule structure that could be easily exploited by our reasoning services (see Section 5.2). Our aim, ultimately, is to adopt SWRL as the representational medium for rule-based reasoning.

4 Technology Infrastructure

The technological infrastructure of the TDS consists of a rich variety of Semantic Web technologies such as OWL, RDF, RDFS and OWL-S, in addition to expert systems technologies such as CLIPS [18]. We also aimed to exploit, where appropriate the technologies delivered as part of the AKT initiative [10, 19]. The AKT initiative is a major UK research initiative aimed at delivering next generation technologies for the Semantic Web. In subsequent sections we present a brief overview of some of the technologies exploited in the context of the current initiative.

4.1 OWL-S

OWL-S [20] provides Web Service providers with a set of language constructs for describing the capabilities and properties of their Web Service portfolios. The appropriate description of Web Services is what makes the vision of the Semantic Web as a set of loosely-coupled, knowledge-oriented services a realistic possibility. OWL-S was used to provide a semantically-rich characterization of the capabilities and characteristics of the semantic web services developed as part of this initiative.

4.2 3Store

The 3Store is an RDF Triplestore implemented on top of a MySQL database [21]. It was developed in the context of the AKT initiative to provide a scalable RDF retrieval and reasoning service. The 3Store is implemented on top of a MySQL database engine, which can be manipulated using conventional queries formulated in SQL. However, in order to provide more sophisticated query capabilities, the 3Store incorporates an RDQL [22] interface. The 3Store RDQL engine transforms an RDQL query into a SQL query, which can

then be executed against the RDBMS representation of the RDF data. A key advantage of 3Store technology as opposed to competing RDF storage and retrieval solutions, such as Jena [6], concerns the speed at which query results can be returned. 3Store can return query results within a few milliseconds, which is a pre-requisite of real-time Semantic Web applications that involve the rapid execution of multiple queries within a limited timeframe.

4.3 CLIPS

CLIPS is an expert system shell that co-opts both a rule-based inference engine with object-oriented programming facilities [19]. CLIPS was used to implement a number of reasoning services each providing reasoning and decision-support capabilities as part of the service portfolio of the system (see Section 5). These services were accessible via a Web Services interface, which enabled easy invocation and exploitation of the reasoning and decision-making processes implemented by the services. Typically, these services queried the 3Store knowledge repository (see Section 4.2) via the knowledge storage services (see Section 5.1) in order to dynamically populate the CLIPS expert system shell environment.

5 Service Portfolio

Much as the World Wide Web benefits from its distributed nature in terms of scalability and robustness, we expect the same to be true of the Semantic Web. One means of delivering scalable and robust solutions to knowledge-related problems in the context of the Semantic Web is to decompose an application's core competencies into relatively independent, loosely-coupled modules. By loosely-coupled we mean that the interactions between system components are not rigidly specified at design time, but rather that existing system components may opportunistically exploit new services as they become available. This vision of the Semantic Web as a system of distributed services that cooperate in a coordinated manner to solve knowledge-oriented problems is at the heart of contemporary research initiatives aimed at exploiting the full potential of the Semantic Web. Typically such initiatives are concerned with one of three challenges facing the widespread adoption of Semantic Web technologies: service description (how should semantic web services be described to enable their opportunistic discovery and exploitation?), service discovery (what is the best means of facilitating the discovery of new services?) and service composition (what is the best way to orchestrate the activities of multiple services in order to achieve a desired knowledge outcome?).

In the context of the current initiative we have decomposed the functionality of the TDS into a set of web services that communicate via HTTP and collaborate to

provide the knowledge and reasoning capabilities demanded of the TDS. Currently, these services are explicitly configured and organized, and are referred to using hard-coded URLs in the relevant system components. We are, however, moving towards a system aimed at addressing the aforementioned problems of service description, discovery and composition, which will enable the new services to be incorporated and make the system less dependent on specific service instances.

Five main types of semantic web service are exposed by the system. These services cooperate to provide the full range of information retrieval knowledge processing and visualization capabilities alluded to in the context of the functional characterization of the system (see Section 2).

5.1 Knowledge Storage Services

These services relate to the persistent storage of knowledge structures in a knowledge repository. Our storage medium is based on 3Store technology (see Section 4.2) [21]. The 3Store provides a Web Services interface to its query engine, which accepts RDF Data Query Language (RDQL) [22] queries and returns query results as RDF documents. RDQL [22] is an SQL-like language for performing sub-graph pattern matches over an RDF graph, returning bindings for the variables specified in its `SELECT` clause. Another set of Web Services enables the population of the knowledge repository. In this case, the knowledge structures are serialized as RDF/XML and passed to the web service interface. The population of the knowledge repository consists in the assertion of instances of the knowledge types specified in the domain ontology. In most cases the population of the knowledge repository with knowledge instances occurs automatically as an adjunct to the operation of the information harvesters (see Section 5.4)

5.2 Reasoning Services

The reasoning services developed as part of the current initiative provide the backbone for many of the decision-making capabilities of the TDS. Such capabilities subsume a number of information-fusion relation activities as well as more general inferential processes. The outcomes of the reasoning services are generally used to modify or refine the conceptual model of the system or to yield decision outcomes that are strategically aligned with the objectives of humanitarian relief efforts.

5.3 Information Dissemination Services

Information dissemination services fulfill the objectives of the system vis-à-vis the active notification of agents who have subscribed to the actual or predicted occurrence of particular events. Such services relate to the system capabilities discussed in Section 2.5. In general, information dissemination services allow external agents to register their interest in the occurrence of particular

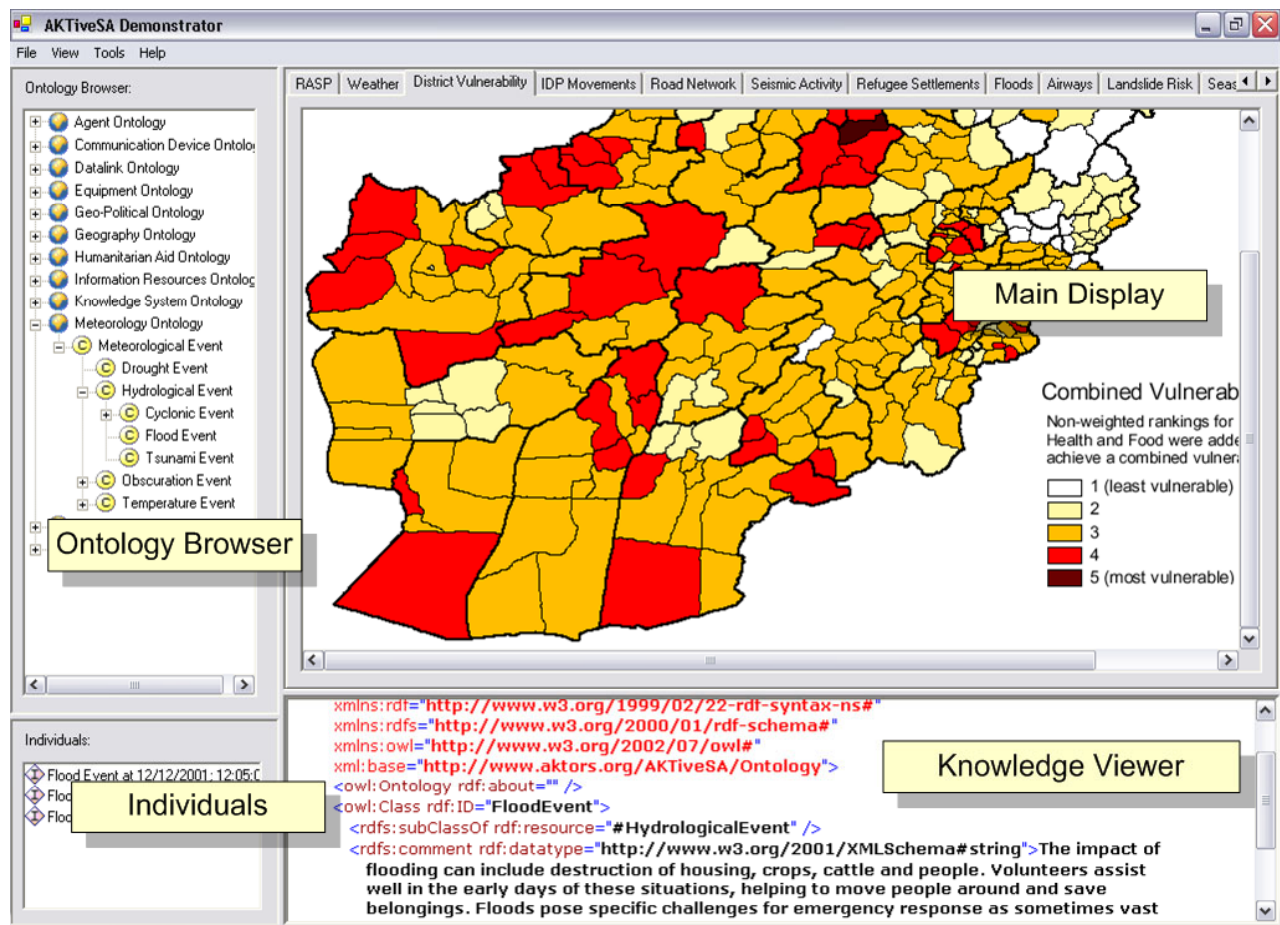


Figure 1. AKTiveSA Demonstrator

types of event, where an event is defined with respect to the ontology developed for the domain.

5.4 Information Retrieval Services

The knowledge and reasoning capabilities of the TDS depend on access to appropriately formatted representations of real-world events. To ensure that the system always has access to the latest information about events in the scenario it is necessary for services to respond to information inputs in an appropriate manner, or, as is more generally the case, to actively harvest information from available sources. Most of our information retrieval services exist as information harvesters that extract information from Web sites, databases, spreadsheets and other sources, convert it into OWL and assert it into the 3Store knowledge repository. A number of models for content acquisition are adopted by the information retrieval services, including the scraping of publicly available data from institutional web sites [23].

5.5 Visualization Services

The purpose of visualization services is to provide a graphical representation of scenario information and

knowledge system output. As was discussed in Section 2.6 the generation of a graphical display is complicated by a number of factors including the cognitive and perceptual profile of an external agent (e.g. human vs. software agent), the organizational affiliation of the agent and their role within the organization (e.g. military vs. civilian personnel) and the characteristics of the display device through which information is to be communicated (e.g. 19 inch color monitor vs. handheld PDA). Although we have selectively focused on services that target the visual modality in the context of the current project it is important to bear in mind that the notion of 'visualization' could be extended to include other modalities. A number of projects within the DIF DTC are investigating the impact of cross-modal information transfer as a means of establishing a more effective distribution of cognitive and perceptual resources under high workload conditions. In some cases the use of auditory information and/or vibrotactile displays may present more effective means of communicating information to human operatives.

6 Visualization and Interaction

Although the nature of the user interface varies according to the type of user accessing the system, we have developed a standard interface as a means of showcasing the capabilities of the TDS to a variety of user groups (see Figure 1). The interface integrates the services described in Section 5 into a single application for monitoring scenario events and knowledge-based system (KBS) activity. The main interface is divided into a number of separate windows or panes, each of which provides information about different aspects of the scenario or the operation of the TDS.

The Main Display window (see Figure 1) provides multiple visualizations of the operational environment with respect to the geo-spatial domain. Each of the tabs within this window provides access to different subsets of information available in the operational environment. For example, the RASP (Recognized Air and Surface Picture) tab depicts the disposition of military assets, while the Weather tab displays information about current weather conditions. The graphical conventions adopted for the display of information in each of the panels can be customized to suit the preferences of particular users via a preferences panel. In addition, the information available from different tabs can be combined in the form of map overlays to provide multi-dimensional compound visualizations of domain information. Each of the tabs features a zoom capability, which enables a user to visualize the information at a variety of spatial scales. While we have provided a fixed set of tabs for the purposes of demonstration it is important to point out that the application features a pluggable component architecture in which new tabs can be developed and incorporated into the application as required.

The information displayed by each tab is relative to a particular time slot in the scenario timeline. In order to provide access to historical information, and to enable easy navigation through the scenario, the AKTiveSA demonstrator provides access to a set of media controls. These controls enable the user to manually set the current time setting for the system, to replay past events at different speeds and to pause the scenario entirely. Such facilities allow users to more closely inspect recent events or to analyze the operation of the reasoning services at different points in the scenario timeline.

Detailed information about any object or event in the scenario can be accessed by selecting the object from one of the tabs provided by the Main Display window. Information about the object is presented in the Knowledge Viewer window (see Figure 1), which displays an ontologically-motivated description of the object based on the information contained in the 3Store knowledge repository (see Section 4.2). The default format for information contained in this window is RDF; however, the user can customize the presentation format

by selecting from one of a number of XSLT stylesheets. New stylesheets can be authored and installed as required by the end user.

The Ontology Browser window (see Figure 1) provides a hierarchical visualization of the knowledge types and instances hosted within the domain ontology developed for the current problem domain. Detailed information about a particular knowledge type or object in the ontology can be accessed by selecting the corresponding node in the treeview control. Information about the selected object is displayed in the Knowledge Viewer window. Note that any individuals, or instances, of the selected knowledge type are enumerated in the Individuals window (see Figure 1). Each of these individuals may have a corresponding graphical representation in the context of the Main Display window. Therefore, the information displayed by the Knowledge Viewer window for a particular individual will be the same irrespective of whether it is selected from the Individuals window or from a tab hosted by the Main Display window.

In addition to the aforementioned windows, the TDS features KBS Output and Information Transfer windows. These windows can be accessed via the View menu option in the main application window (see Figure 1). The KBS output window provides information about the processing activity of the various reasoning services (see Section 5.2) enlisted for the application. The content of this window is important in terms of analyzing the reasoning and decision-making capabilities of the system at different points in the scenario timeline. It provides a means for the system to justify decision outcomes with respect to a natural language trace of reasoning activity. Such information is invaluable in terms of validating and debugging the knowledge and reasoning capabilities of the system. The Information Transfer window details the progress made with respect to the receipt or retrieval of information from various sources. The window provides a natural language trace of the processing activities of the information retrieval services (see Section 5.4).

7 Conclusion

This paper describes our progress with respect to the development of a Technical Demonstrator System (TDS) for improved situational awareness in the domain of humanitarian operations, especially when such operations occur against a backdrop of ongoing military conflict. We have demonstrated how information, harvested from a variety of physically disparate and semantically heterogeneous data sources, can be interpreted with respect to formal ontological characterizations of the problem domain in order to establish a common repository of semantically-rich, conceptual-level representations of real-world events and objects. We have additionally shown how such representations can be

exploited by knowledge and reasoning services to provide useful decision outcomes that are strategically aligned with the goals and objectives of both military and humanitarian agencies. Knowledge services, operating in the context of the Semantic Web, can assist with a variety of fusion-relation activities that are necessary for improved situational awareness in this area. In addition, such services can be used to filter unwanted or irrelevant information from the input stream and deliver report streams that are specifically geared to the information processing objectives and cognitive profile of distinct information consumers (including the TDS itself). Such techniques help to avoid situations of information overload that might otherwise result from fusion-related processes and the exploitation of vast information repositories, such as the World Wide Web.

Acronyms & Abbreviations

AKT	Advanced Knowledge Technologies
API	Application Programmatic Interface
BBC	British Broadcasting Association
CLIPS	C Language Integrated Production System
DIG	DL Implementation Group
DL	Description Logic
HTTP	Hyper Text Transfer Protocol
KBS	Knowledge-Based System
MoD	Ministry of Defence
OWL	Web Ontology Language (W3C)
OWL-S	OWL-based Web Services Ontology
PCPACK	PC-Enabled Portable ACquisition of Knowledge
PDA	Personal Digital Assistant
RDBMS	Relational Database Management System
RDF	Resource Description Framework
RDFS	RDF Vocabulary Description Language
RDQL	RDF Data Query Language
SQL	Structured Query Language
SWRL	Semantic Web Rule Language
SWS	Semantic Web Services
TDS	Technical Demonstrator System
UK	United Kingdom
URL	Uniform Resource Locator
US	United States
W3C	World Wide Web Consortium
XML	eXtensible Markup Language
XSLT	eXtensible Stylesheet Language Transformations

References

- [1] <http://www.w3.org/2001/sw/>, W3C Semantic Web Website.
- [2] G. Antoniou and F. van Harmelen, "Web Ontology Language: OWL," in *Handbook on Ontologies*, S. Steffen and R. Studer, Eds. Berlin: Springer Verlag, 2003.
- [3] D. L. McGuinness and F. van Harmelen, "OWL Web Ontology Language Overview," 2004.
- [4] K. Holger, R. W. Ferguson, N. F. Noy, and M. A. Musen, "The Protege OWL Plugin: An Open Development Environment for Semantic Web Applications," presented at The Third International Semantic Web Conference - ISWC 2004, Hiroshima, Japan, 2004.
- [5] N. F. Noy, M. Sintek, S. Decker, M. Crubezy, R. W. Ferguson, and M. A. Musen, "Creating Semantic Web Contents with Protege-2000," *IEEE Intelligent Systems*, vol. 16, pp. 60-71, 2001.
- [6] <http://jena.sourceforge.net/>, Jena Semantic Web Framework Homepage.
- [7] R. Stevens, R. McEntire, C. M. Goble, M. Greenwood, J. Zhao, A. Wipat, and P. Li, "myGrid and the drug discovery process," *Drug Discovery Today: BIOSILICO*, vol. 2, pp. 140-148, 2004.
- [8] <http://www.mygrid.org.uk/>, myGrid Website.
- [9] N. R. Shadbolt, F. Ciravegna, J. Domingue, W. Hall, E. Motta, K. O'Hara, D. Robertson, D. Sleeman, A. Tate, and Y. Wilks, "Advanced Knowledge Technologies at the Midterm: Tools and Methods for the Semantic Web," in *Advanced Knowledge Technologies: Selected Papers*, N. R. Shadbolt and K. O'Hara, Eds. Southampton: AKT, University of Southampton, 2004.
- [10] <http://www.aktors.org/akt/>, Advanced Knowledge Technologies Website.
- [11] <http://www.sts.tu-harburg.de/~r.f.moeller/racer/>, RACER Website.
- [12] N. R. Shadbolt and M. Burton, "Knowledge elicitation: a systematic approach," in *Evaluation of Human Work: A Practical Ergonomics Methodology*, J. R. Wilson and E. N. Corlett, Eds., 2nd ed. London, England: Taylor and Francis, 1990, pp. 406-440.
- [13] N. R. Shadbolt, K. O'Hara, and L. Crow, "The Experimental Evaluation of Knowledge Acquisition Techniques and Methods: History, Problems and New Directions," *International Journal of Human-Computer Studies*, vol. 51, pp. 729-755, 1999.
- [14] <http://www.epistemics.co.uk/>, Epistemics Website.
- [15] G. Schreiber, H. Akkermans, A. Anjewierden, R. de Hoog, N. R. Shadbolt, W. Van de Velde, and B. Weilinga, *Knowledge Engineering and Management: The CommonKADS Methodology*. Massachusetts: MIT Press, 2000.
- [16] I. Horrocks, P. F. Patel-Schneider, H. Boley, S. Tabet, G. Benjamin, and M. Dean, "SWRL: A Semantic Web Rule Language Combining OWL and RuleML," World Wide Web Consortium, 2004, <http://www.w3.org/Submission/SWRL/>.
- [17] <http://www.ruleml.org/>, RuleML Homepage.
- [18] "CLIPS Expert System Shell Website."
- [19] <http://www.ghg.net/clips/CLIPS.html>, CLIPS Expert System Shell.
- [20] D. Martin, M. Burstein, J. Hobbs, O. Lassila, D. McDermott, S. McIlraith, S. Narayanan, M. Paulucci, B. Parsia, T. Payne, S. Evren, S. Naveen, and K. Sycara, "OWL-S: Semantic Markup for Web Services," World Wide Web Consortium, 2004, <http://www.w3.org/Submission/2004/SUBM-OWL-S-20041122/>.
- [21] S. Harris and N. Gibbins, "3store: Efficient Bulk RDF Storage," presented at The 1st International Workshop on Practical and Scalable Semantic Web Systems, Sanibel Island, Florida, USA, 2003.
- [22] A. Seaborne, "RDQL - A Query Language for RDF," World Wide Web Consortium, 2004, <http://www.w3.org/Submission/2004/SUBM-RDQL-20040109/>.
- [23] T. Leonard and H. Glaser, "Large-Scale Acquisition and Maintenance from the Web without Source Access," presented at The 4th Workshop on Knowledge Markup and Semantic Annotation, 2001.