UNCLASSIFIED

Knowledge-Intensive Fusion for Situational Awareness

Technical Progress Report

28 December 2004

Reference DTC/WP500/Progress

Version 1

Prepared by ....Paul Smart

Reviewed by.. Prof. Nigel Shadbolt

The investigation that is the subject of this report was initiated by the University of Southampton under the terms of Contract 8.14 for Prof. Nigel R. Shadbolt, Professor of Computer Science, University of Southampton.

ALL RECIPIENTS OF THIS REPORT ARE ADVISED THAT IT MUST NOT BE COPIED IN WHOLE OR IN PART OR BE GIVEN FURTHER DISTRIBUTION OUTSIDE THE AUTHORITY OF THE PRIME CONTRACTOR WITHOUT THE WRITTEN APPROVAL OF THE DIRECTOR OF COMMUNICATIONS INFORMATION AND SIGNAL PROCESSING

For
Prof. Nigel R. Shadbolt
Professor of Computer Science
University of Southampton
Electronics and Computer Science
University of Southampton
Southampton
SO17 1BJ

Distribution

1 copy..........................University of Southampton

Acknowledgements

This work (project 8.14) was supported by the Data and Information Fusion Defence Technology Centre (DIF DTC), a consortium of academic and industrial partners headed by General Dynamics UK Ltd. The DIF DTC was established by the UK Ministry of Defence (MOD) to undertake leading-edge research in the area of data and information fusion (http://www.difdtc.com).

Amendments

Version	Date of Issue	Change Reference	Remarks
1	28 December 2004		

Executive Summary

This document details the technical progress attained in the context of DTC Project 8.14. It reviews the work undertaken to date and specifies the work to be undertaken in the remainder of the project. Issues such as the capabilities of the prospective system as well as its technological and ontological infrastructure are discussed, as are issues of project risk, system evaluation, opportunities for collaboration, routes of exploitation and project deliverables. In essence the document provides a snapshot of the current status of the project and lays the foundation for further technical progress. The primary purpose of the document is to establish a consensus and common purpose with respect to the goals and objectives of the project and document any outstanding concerns about project progress. This document is not part of the contracted deliverables required by the funding authority, but could be submitted as a means of eliciting feedback about the project and ensuring stakeholder confidence in the project's progress.

Contents

1 Introduction

1.1 Project Background

This project (contract no: 8.14) addresses the development of a Technical Demonstrator System (TDS) to showcase the ability of knowledge technologies to improve situational awareness via intelligent information fusion. The work is being undertaken by the University of Southampton[1] as part of the MOD's DIF DTC initiative[2], which represents a formal collaborative agreement between industry and academic experts to generate and enhance the defensive capabilities of UK military forces.

This project is an extension to an earlier initiative, called FloodSim[3], which demonstrated how semantically-enriched information, interpreted against the backdrop of formal ontologies, could be used to improve situational awareness with respect to humanitarian relief operations. The current project has a similar objective in that it aims to improve operational effectiveness in the planning, coordination and delivery of humanitarian relief operations by improving the situational awareness of executive decision makers. At the heart of the current project is the need to receive, and in some cases actively acquire, information that can be subsequently assimilated (fused) into a coherent representation of the current operational environment. This operational environment provides a framework within which information can be disseminated to executive agencies in a manner that befits their specific epistemic and representational requirements (i.e. the system supports multifarious visualizations of the underlying situation in a manner that augments the problem-solving competency of key knowledge workers). It also provides a basis for reasoning activities aimed at improving the operational effectiveness of decisions undertaken in regard to humanitarian initiatives, e.g. alerting operatives to information that may have been overlooked, advising as to the best course of action, selectively presenting information of immediate strategic relevance (thereby avoiding the notorious problem of information overload), etc. In general, there are 4 major objectives for this project:

1. to leverage increased operational effectiveness from improved situational awareness in the planning, coordination and delivery of humanitarian relief efforts
2. to demonstrate the effective use and exploitation of ontological characterizations of the target domain in the interpretation, integration and presentation of semantically-enriched information
3. to exploit knowledge-rich contingencies supporting competent performance in the target domain in order to yield operationally effective decision outcomes that are strategically aligned with the goals and objectives of humanitarian agencies or those involved in relief efforts

[1] http://www.ecs.soton.ac.uk/

[2] http://www.difdtc.com/

[3] http://www.aktors.org/technologies/floodsim/

4. to exploit the technologies and methods developed in the context of the AKT (Advanced Knowledge Technologies) initiative[4] as a means of highlighting the value and utility of these techniques for a variety of knowledge-intensive applications

These aforementioned objectives are reflected in the capabilities (see Section 3) of the proposed TDS, which represents the key technological artefact to be delivered by this project.

1.2 Document Purpose

The aim of the current document is to outline the current status of Project 8.14 in respect of its key objectives (see Section 1.1). The document also aims to detail the work to be undertaken in the remainder of the project, the project milestones to be reached with respect to this work, and the deliverables that will result from the successful completion of these milestones. The current document therefore provides a snapshot of technical progress currently attained in the project and formally specifies the technical progress to be achieved in the remainder of the project.

Since the current document aims to provide a comprehensive vision of the functionality of the prospective system, and the work to be undertaken in the technical realization of this vision, it is important that the views expressed in this document are made explicit and communicated to project stakeholders as a means of establishing consensus with respect to the technical objectives, requirements and commitments of the project.

Finally, although this document is not part of the formal list of deliverables to be submitted to General Dynamics UK Ltd[5], as part of the contractual obligations of the current project, nothing in this document prohibits the submission of this document as a formal deliverable if required. Indeed, submission of this document may be beneficial from the perspective of improving stakeholder confidence and ensuring that the perceived functionality of the proposed system is consistent with the customer's requirements and expectations. Critical feedback is welcomed from all parties as a means of ensuring the technical excellence and relevance of the project to the UK's defence requirements.

1.3 Document Scope

The scope of this document is limited to a discussion of the technical progress thus far attained in the project and a review of the technical progress yet to be achieved. The current document should not necessarily be construed as a fixed mandate for project actions. Rather, it should be regarded as the most detailed functional characterization of the prospective TDS at the current point in the project timeline. In particular, the current document should not limit or commit project stakeholders to a particular implementation strategy or procedure for the realization of the technical and intellectual goals discussed herein.

---

[4] http://www.aktors.org/akt/

[5] http://www.generaldynamics.uk.com/

1.4 Document Basis

This document builds on a previous document[6] submitted to General Dynamics UK Ltd. as part of the delivery requirements of the current project. The former document, entitled 'Scenario Specification', provided a provisional specification of a scenario involving a variety of humanitarian relief actions occurring against a backdrop of ongoing military conflict. In addition to the scenario specification, this former document also detailed a number of project actions that should be implemented following document submission. The current report refines and updates these actions and comments on the progress attained with respect to their implementation.

1.5 Document Structure

The structure of the current document reflects the document aims detailed in Section 1.2. Section 2 provides an overview of the current status of the project and the significant milestones achieved thus far. Section 3 describes the capabilities of the prospective TDS especially with respect to its knowledge processing and information fusion activities. As the major technological artefact to be delivered by this project a detailed functional characterization of the prospective system is a central aspect of the requirements specification activity. Section 4 provides an overview of the range of technologies likely to be exploited in the context of the current development initiative. Since the aims of the current project are closely aligned with ongoing research initiatives in the IAM[7] group at the University of Southampton, e.g. AKT[8], some of the technologies developed in the context of those projects will be exploited in the context of the current initiative, e.g. 3Store technology[9]. Furthermore, a range of extant technologies are likely to prove useful both at the modelling and implementation levels of the prospective system. These include, but are not necessarily limited to, the UML, the CommonKADS methodology, eKADS, PCPACK, OWL, RDF and JESS. Section 5 describes, in summary form, the knowledge infrastructure of the target domain. The knowledge infrastructure subsumes all those knowledge structures, e.g. domain conceptualizations, problem-solving methods, knowledge-rich contingencies, etc., that play a role in terms of ensuring problem-solving success in the target domain, i.e. in the area of humanitarian relief operations. The specification of the knowledge infrastructure should be described at the 'knowledge level' (Newell, 1980) and it is therefore necessarily independent of implementation detail, even to the extent of eschewing the representational biases of particular knowledge modelling/representation languages. Although it is perfectly acceptable to countenance one particular set of representational formalisms for the express purpose of implementing intelligent software, based on the software's functional requirements and the representational leverage afforded by different implementation mechanisms, the same formalisms are seldom required for initial knowledge modelling activities in which the emphasis is on stakeholder communication and knowledge validation. The knowledge models developed for the current initiative therefore focus on representational techniques

---

6 see 'Scenario Specification' (document reference: DTC/WP100/Scenario).

7 http://devel.iam.ecs.soton.ac.uk/

8 http://www.aktors.org/

9 http://triplestore.aktors.org/

that easily promote a shared understanding and communication of the essential knowledge structures used in the domain. Section 6 focuses on system evaluation, specifically the issue of whether the system adequately realizes its technological and intellectual objectives. Section 7 discusses potential collaborative opportunities within the next round of DTC funding. The current project is currently supported by a grant awarded to the University of Southampton as part of the DTC's Phase I initiative. The Phase II initiative will commence in February 2006 where the emphasis will be on perceived capability gaps in the MOD's defensive capabilities, as well as further development of technologies delivered in the context of the Phase I projects. In the selection of Phase II projects, preference may be given to project proposals that emphasize collaboration between the Phase I research groups, especially to the extent that such proposals aim to co-opt and exploit the results of multiple Phase I initiatives. The identification of collaborative opportunities depends on a detailed understanding of the activities of other research groups within the DTC, and to this end Section 7 provides a concise overview of all projects and project participants within the DIF DTC. Section 8 describes the potential exploitation routes for the technologies developed in the context of the current initiative. Section 9 details the deliverable items to be provided by the current project. Some of these items correspond to the formal deliverables to be delivered to the funding authority as part if the contractual arrangements specified for the project. However, a number of other deliverables may be delivered en route to the successful resolution of the project. These include research papers, knowledge models, software artefacts, progress reports, project plans and project management materials, as well as largely intangible products such as intellectual capital. Finally, Section 10 addresses a number of risks and uncertainties surrounding the current project in the form of a feasibility assessment. Issues raised in this section could constitute the basis for discussion at future stakeholder meetings.

As a description of a knowledge-rich application domain, namely the domain of military operations and humanitarian relief efforts, this document uses a number of acronyms and abbreviations. These are detailed in Appendix A. Appendix B lists the references of source materials and documents used in the compilation of this report.

2 Project Status

This section describes the current status of the project with respect to the project objectives outlined in Section 1.1.

2.1 Scenario Specification

A scenario specification document (reference: DTC/WP100/Scenario) was submitted to the relevant funding authority in April 2004 as part of the contractual commitments of the current project. The aim of this document was to detail a scenario to showcase the capabilities of the proposed TDS in a number of capability areas (see Section 3). As such, a number of constraints were developed to delimit the space of possible candidate scenarios:

1. **Realism**: the scenario must be sufficiently realistic, i.e. aligned with the same conditions and constraints that govern real-world cases of humanitarian intervention.
2. **Requirement for Humanitarian Intervention**: Since the primary focus of the technical demonstrator is geared towards processing information relevant to humanitarian relief efforts, the scenario must clearly provide a requirement for humanitarian intervention.
3. **Background Military Conflict**: the humanitarian operations featured in the scenario must occur against a backdrop of ongoing military conflict
4. **Dynamic Tactical Picture**: the scenario must be of sufficient complexity and richness to tax the information processing capabilities of the prospective system to its full extent.
5. **Noisy Information**: the system must be sufficiently robust to cope with noisy and conflicting information. This will test the ability of the system to filter information based on its credibility, trustworthiness, timeliness, consistency and other evaluative criteria.

Based on these constraints a scenario was developed that described both acute and chronic humanitarian events occurring against a backdrop of ongoing military conflict[10].

The key purpose of the scenario specification exercise was to facilitate the programme of knowledge engineering to be undertaken in regard to the current initiative. In particular, a scenario specification can assist with the following knowledge engineering activities[11]:

- **Knowledge Capture**: the acquisition of domain knowledge, typically from human subject matter experts.

[10] The validity of this scenario was questioned at a recent DTC review meeting (see Section 2.7). The primary concern related to the availability of subject matter experts for knowledge elicitation. This issue is addressed in Section 2.7.

[11] See the aforementioned scenario specification document for further information.

- **Validation**: the validation of acquired knowledge, e.g. do the observed decision outcomes match the required decision outcomes.
- **Testing**: experimental testing of the system to determine its robustness and integrity in the face of a variety of challenges and manipulations.
- **Demonstration**: the adequate communication of the system's capabilities to project stakeholders and other interested parties.
- **Evaluation**: the assessment of the system's performance with respect to a number of measures of operational effectiveness (MOEs). The evaluation of the system with respect to these measures is important in terms of determining whether the system has lived up to expectations regarding its putative ability to improve situational awareness and operational effectiveness.

While the scenario specification initiative has proved useful in a number of knowledge engineering activities, e.g. modelling the knowledge infrastructure of the target domain, the need for further development and enrichment of the scenario should not be overlooked. In particular, it is of utmost importance that the scenario should be characterized not only in terms of the information objects that represent aspects of the scenario, but also in terms of the messages that communicate information about the scenario to executive agencies with varying degrees of fidelity. As such we recommend that subsequent development efforts should focus on the detailed instantiation of the scenario in accordance with the aforementioned constraints and that these refinements should serve as the basis for TDS demonstration. One strategy of potential value in this respect is to characterize the scenario in terms of a number of data objects and then to characterize each object at each point in the scenario timeline. Since this is a rather labour-intensive process it is possible that a dedicated scenario editor facility could be developed to assist with this process.

2.2 Subject Matter Experts: Identification and Contact

Contact with subject matter experts is important for any knowledge engineering initiative. Not only are such individuals able to comment on the integrity and validity of acquired knowledge, they may also possess information that is impossible to acquire from other information resources, e.g. tacit procedural knowledge. Experts can also provide valuable feedback about the kind of tasks in which the proposed system should engage and the manner in which the system should inter-operate with a variety of other agents.

A number of efforts have been made to contact humanitarian agencies with respect to the identification and recruitment of human subject matter experts that could assist with the functional characterization of the TDS and knowledge capture initiatives. These efforts have entailed contacting humanitarian agencies, providing a brief description of the project, and enquiring as to whether help could be given. At the present time such efforts have failed to secure the cooperation of any humanitarian agencies; however, we aim to continue with this effort throughout the project lifecycle[12].

[12] One possible contact was communicated to Nigel Shadbolt, namely the wife of Tim Thorpe. This contact should be pursued as soon as possible.

2.3 Knowledge Modelling

A partial knowledge model of the problem domain was submitted in October 2004 as the basis for formal ontological characterizations of the knowledge infrastructure surrounding humanitarian relief efforts. This model divided the domain into a number of distinct areas or domain schemas (see Section 5.4), which both simplifies the structure of the knowledge model and promotes reuse of existing knowledge model components. Each of the domain schemas corresponds to a distinct ontology in the knowledge repository of the proposed system. The aim, in terms of the knowledge engineering initiative, is to provide knowledge model specifications for each of these ontology areas using the Knowledge Model component of the CommonKADS model suite (Schreiber et al, 2000). The purpose of the CommonKADS Knowledge Model is to:

> "explicate in detail the types and structures of the knowledge used in performing a task. It provides an implementation-independent description of the role that different knowledge components play in problem-solving, in a way that is understandable for humans. This makes the knowledge model *an important vehicle for communication* with experts and users about the problem-solving aspects of a knowledge system, during both development and system execution." (Schreiber et al, 2000; pg 19) [my emphasis]

The key point here is that the model should serve as a vehicle for communication about the knowledge infrastructure of the domain between individuals who may have very different interests and roles in relation to the project, e.g. SMEs, project managers, knowledge engineers, system developers, and so forth. Accordingly, we advocate a two-phase approach to the development of the knowledge system component of the TDS[13]: in the first phase knowledge models are constructed (using the CommonKADS Knowledge Model) to reflect the key aspects of the static knowledge infrastructure and problem-solving process in a manner that is independent of any implementation-specific detail or representational concern; in the second phase these models are translated into an implementation-specific knowledge representation language that supports the kinds of semantic expressivity and inference logic required of the prospective system[14]. For the most part we expect that the transition from Phase I to Phase II models can be accomplished automatically (see Section 2.4); however, there are some caveats here. In particular, OWL[15] does not provide support for the kinds of rule contingencies deemed necessary for the cognitive profile of the prospective system. The CommonKADS methodology adopts a representational strategy within which this type of knowledge is represented in a production rule-like format, similar to conventional expert systems based on production systems technology, e.g. CLIPS. OWL does not yet avail itself of such contingencies and as a result only supports a limited form of inferential capability,

---

[13] In practice, the knowledge modelling process often assumes the form of a three-step process. An initial phase of knowledge acquisition and informal modelling typically precedes the provision of formal or semi-formal knowledge models. This tripartite system of model formality is reflected in the use of different tools for each stage of the knowledge modelling process. In the initial stage the tendency is to use tools that are specifically designed to support the knowledge capture process, e.g. PCPACK, while in the latter stages, where the emphasis is on formal modelling, other types of technology come to the fore, e.g. eKADS and Protégé.

[14] In all likelihood this final delivery format for domain knowledge will consist of a series of OWL-based ontologies.

[15] http://www.w3.org/2001/sw/WebOnt/

namely one supported by the logical structure of the conceptual model, e.g. if 'brother-of(john, bill)' and 'brother-of' is a transitive relationship, then it can be inferred that 'brother-of(bill, john)'[16]. To circumvent the limitations imposed by this rather restrictive model of logico-deductive inference, we anticipate three possible options for the current project[17]:

1. extension of OWL to include the required representational formalisms,
2. inclusion of OWL-independent knowledge representations that are integrated with OWL constructs at the application level,
3. the development of meta-representational formalisms (i.e. representations of rule contingencies) built on top of the existing OWL constructs that will support the requisite degree of representational richness and inferential complexity.

Of these three strategies, the last strategy seems the most favourable to us in the context of the current development initiative and is the one that will be pursued with the greatest enthusiasm. Since this approach is not currently widespread in the knowledge engineering community we anticipate that these efforts could yield an important intellectual outcome worthy of communication in the academic and commercial communities.

2.4 Technology Development

The main emphasis of work undertaken since the submission of the scenario document has been to capture relevant domain knowledge and construct knowledge models that both reflect the conceptual infrastructure of the problem domain and that support the kinds of reasoning and decision-making capabilities expected of the final system (see Section 3). The key technology used in this respect has been the PCPACK toolkit[18] marketed by Epistemics[19]. PCPACK consists of a collection of specialized KA tools that serve as the software counterparts to a number of knowledge capture techniques. The tool suite includes the following knowledge capture tools:

- **Laddering Tool**: used to construct hierarchical decompositions of domain conceptualizations, e.g. taxonomic and componential hierarchies
- **Matrix Tool**: used to edit relationships and characterize knowledge objects using a simple 2-dimensional grid

---

[16] OWL additionally supports a form of reasoning called subsumption reasoning in which a reasoning agent can compute the taxonomic (subsumption) hierarchy for a domain based on a characterization of existing domain conceptualizations according to their properties and inter-relationships with other concepts. For example, if we assert that a 'dog' is a type of mammal and has exactly 4 legs and we also assert that a '4-legged-mammal' is a type of mammal and has exactly 4 legs, then a reasoner, such as RACER (see http://www.racer-systems.com/index.phtml) can infer that a 'dog' is also a type of '4-legged-mammal'.

[17] Efforts are currently underway in the W3C community to address this issue via the development of a new language called SWRL (Semantic Web Rule Language). SWRL is a combination of OWL and a dedicated rule implementation language called RuleML. Further information about SWRL can be found on the W3C website (http://www.w3.org/Submission/SWRL/).

[18] http://www.epistemics.co.uk/Notes/55-0-0.htm

[19] A further development effort currently underway at Epistemics aims to deliver a more extensive suite of tools within a web-enabled environment (see http://www.i-kew.com).

- **Protocol Editor Tool**: used to acquire knowledge from source materials using customized marker pens
- **Annotation Tool**: used to associate multimedia and HTML content with any knowledge object
- **Diagramming Tool**: used to illustrate modelling concepts using custom diagrammatic notations, e.g. UML classes
- **Diagram Template Editor Tool**: used to create custom diagrammatic notations to represent key modelling concepts

A common problem encountered when modelling knowledge is the method used to exchange information between knowledge tools, each of which is suited to a particular phase of the knowledge engineering process. This issue was briefly mentioned in Section 2.3 in the context of automatic translations between CommonKADS-based and OWL-based models. The problem is compounded by the fact that different stages of the ontology development process place different emphasis on different types of information and therefore have somewhat different representational concerns. In the context of PCPACK the issue of knowledge model migration is addressed by writing translators that convert the native XML format adopted by PCPACK into the required representational language[20]. To accomplish the automatic translation of PCPACK-based representations into OWL, a specific PCPACK extension was developed courtesy of Mr Steve Swallow[21] at Epistemics. This tool was developed in the context of a larger initiative currently being undertaken by Epistemics: the CO-ODE initiative[22]. This initiative aims to improve the interaction between PCPACK and other knowledge editors such as Protégé.

2.5 Requirements Specification

An initial functional characterization of the proposed system has been provided based on the original project proposal and the current state-of-the-art in knowledge/ontology engineering, the semantic web and HMI design. The aim is to push the boundaries of the what is currently feasible in terms of extant technology and techniques without reneging on the commitment to produce operational end-products that can be adapted for real world use. The current document describes the functional characterization of the system in relation to five capability areas. These capabilities are discussed in greater detail in Section 3.

[20] Unfortunately, given the often iterative process of knowledge engineering, it is often necessary to implement two-way translations in which knowledge constructs that are migrated from one format to another are, at some later time, required to be re-converted to their original format. The problems posed by the two-way translation requirement are beyond the scope of the current project and are not considered further.

[21] email: steve.swallow@epistemics.co.uk

[22] http://www.co-ode.org/

2.6 Staff Recruitment and Training

Paul Smart, the primary author of this report, has recently been recruited to serve as principal knowledge engineer on the current project. A further research post is expected to be filled in due course. Since his induction on 1st November 2004, Paul has undertaken a number of training initiatives including attendance at an RDF TripleStore seminar organized by Hugh Glaser. Paul has also spent time becoming familiar with the technical requirements of the project and addressing any perceived capability gaps.

2.7 Presentations, Meetings and Other Communications

Since the submission of the scenario specification report to the funding authority two meetings have taken place: one focused on a review of technical progress in the project and was attended by members of General Dynamics UK Ltd., the other was a DTC Theme Meeting organized by Martin Ferry who is head of the DIF DTC theme entitled 'Situational Awareness and Human Factors'. Detailed notes regarding the latter meeting were distributed to project stakeholders in the form of a document entitled 'DTC Theme Meeting' (reference: DTC/Notes-28-10-2004#1). In terms of the former technical review meeting a number of issues were raised by the reviewers including decisions surrounding the choice of scenario and relevant contact with other DTC projects. In terms of the latter concern the reviewers were keen to emphasize that contact with Project 7.6 and projects within the Agents and Architectures theme could be relevant for subsequent development efforts. Other contacts of potential value include Nick Beswick (in Andy Tilbrook's team at General Dynamics) and Panos Louvieris at the University of Surrey. These contacts were mentioned in relation to the visualization technologies that could be exploited in the relation to the current project. The DTC review team also commented on the choice of location for the scenario as detailed in the scenario specification document. A suggestion was made to switch the scenario location from Afghanistan to Cambodia because humanitarian operations are currently winding up in Cambodia and a greater number of SMEs may therefore be available with relatively recent operational experience. There are three core concerns we have in relation to this suggestion and which we believe vitiate the recommended alteration of the scenario:

1. It is important that the scenario details events of both military and humanitarian relevance and that the system is able to demonstrate the effective inter-operation of humanitarian and military agencies (see Section 2.1). To our knowledge the humanitarian situation in Cambodia does not entail concurrent military activities involving British military forces.
2. It is important to avoid the capture of situation-specific knowledge in favour of generic rules, conceptualizations and procedures that are common to multiple instances of humanitarian intervention. The decision to exclusively focus the knowledge analysis on a particular instance of humanitarian aid runs the risk of capturing highly specific and detailed knowledge that has little or no relevance to other cases of humanitarian aid.
3. The recommendation to switch the location of the scenario is based on the assumption that the availability of humanitarian aid personnel from one location necessitates an alteration of the scenario to match the experiences of these personnel. Instead, we believe that the experience and expertise of such personnel

will be valuable irrespective of the actual location of the aforementioned scenario. Indeed it will of interest to determine if the knowledge gleaned from one set of experiences can be flexibly applied to a relatively novel situation with subtly different humanitarian aid challenges.

3 System Capability

3.1 Overview

The primary objective of the TDS is to highlight how real-time information, harvested from a rich variety of physically disparate and semantically heterogeneous information sources, can be interpreted with respect to a rich corpus of background knowledge in order to improve the problem-solving competency and operational efficiency of a variety of inter-operating agents. These capability and performance improvements are expected to ride on the profound improvement in situational awareness delivered by the aforementioned system. In particular, the system should avail itself of a number of core information processing capabilities in order to improve situational awareness and operational effectiveness in regard to the planning and implementation of humanitarian aid programmes. This section reviews stakeholder expectations regarding the prospective system's capability in five capability areas:

- **Information Retrieval**: the ability to receive, monitor and actively acquire information of perceived relevance to the current tactical picture.
- **Information Fusion**: the ability 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 picture of the immediate operational environment.
- **Knowledge Processing**: the ability to reason over incoming information streams in order to realize operationally-useful decision outcomes.
- **Information Dissemination**: the ability to identify agents that have registered an interest in the receipt of particular types of information or that require immediate information about some event in the tactical picture based on a characterization of their operational role. This capability area includes the ability to selectively filter incoming information streams along a variety of semiometric dimensions in order to avoid situations of information overload, an ability otherwise known as information triage.
- **Agent Interaction and Visualization**: the ability to support 'effective' visualizations and interact with the user in a manner that improves situational awareness and augments the problem-solving profile of end-user agents and inter-operating systems.

Each capability area is described in more detail in subsequent sections.

3.2 Information Retrieval

The system should be able to source information from diverse sources, e.g. web pages, online databases, tactical datalinks, email notifications, etc., and interpret this information against the backdrop of formal ontological characterizations of the problem domain. Some of these sources may be hardwired into the system, but in other cases the system may be required to actively discover new services and information feeds that are relevant to ongoing tasks, especially since the set of relevant information sources is likely to change during the course of a humanitarian relief effort.

In some cases the TDS will be passive with respect to the receipt of information and in other cases it will need to actively query information sources based on its own knowledge and reasoning requirements (i.e. its 'epistemic hunger'). For example, the system may seek supporting information about the occurrence of an event in order to increase its confidence that the event has actually occurred. In general, similar information received from a variety of information agencies (some of which may operate from different ideological standpoints) will increase the systems confidence in the accuracy of information about events.

Once information has been interpreted, the system should be able to assess the semantic relevance of this information with respect to the goals and objectives of ongoing tasks. Irrelevant information should be filtered from the incoming information streams. This assessment of semantic relevance is similar to an attentional process that selectively focuses on information that is germane to current interests and concerns, i.e. task objectives. As such the cognitive psychology literature on attention may be of relevance to the capabilities of the prospective system in this area.

3.3 Information Fusion

Filtered information should be integrated (fused) into a coherent representation of the current operational environment (hereafter referred to as the tactical picture). The tactical picture provides a God's eye view of the information environment relevant to a task in which the TDS is engaged. This fusion capability is central to the functional integrity of the TDS, but is complicated by the fact that the tactical picture often needs to be actively constructed as opposed to passively received. This notion, which has its analogues in the human perception psychology literature, emphasizes the problems raised by noise in the immediate perceptual environment. For example, how should the system deal with conflicting information from different sources (conflict resolution), how should it update the current tactical picture with respect to new information, especially when that information comes from multiple information sources, and how confident can the system be that the information it receives is an adequate reflection of real-world events and objects. The key to all these problems is the notion of trust, i.e. what level of trust should the system invest in each of the various information streams and inputs it receives. The problem of trust in turn introduces us to the related problems of uncertainty and belief. In particular, the manner in which the system resolves fusion-related issues turns on:

1. the provision of semantically-enriched characterizations of the information sources from which information derives and
2. the maintenance of a belief system that is informed by the provenance and trustworthiness of information sources as specified in the aforementioned semantically-enriched characterizations and
3. the ability to make sensible inferences and executive decisions against the backdrop of a system of aforementioned beliefs

Issues of uncertainty are important here because they have a rather direct impact on the decision-making capability of the system and the resulting confidence the system has in decision outcomes. These may in turn influence the way in which different response options are sanctioned by the TDS.

3.4 Knowledge Processing

Following the formulation of a coherent tactical picture, the system may be required to perform additional reasoning (beyond that discussed in relation to the aforementioned capabilities) in order to infer new information or make decisions relevant to the planning, coordination and delivery of humanitarian operations. In some cases missing or unreliable information may be asserted with a higher degree of confidence as the result of some inferential process undertaken by the TDS. For example, if the TDS can infer that a number of information reports all refer to the same entity, e.g. a humanitarian aid convoy, then the system can consolidate information contained in the reports in order to generate a more stable impression of the entity's characteristics, e.g. its location, size, intended destination, etc. Similarly, the system may be able to infer the scale of a humanitarian disaster based on information about the occurrence of an event (e.g. an earthquake), the magnitude of an 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 certain magnitude may cause structural collapse or tsunami events depending on its epicentral location), and the population density of regions directly affected by the event or its causal precursors or sequelae. Inferential processes aimed at estimating the scale of a humanitarian disaster may, in turn, have a profound impact on subsequent knowledge-based processes, e.g. which subscribers or agencies should be contacted or alerted with respect to the occurrence or impending occurrence of a humanitarian disaster event, what types of humanitarian action should be initiated in the face of a disaster and what types of humanitarian aid resource should be marshalled to pre-empt further suffering or loss of life[23]. Of course confidence limits apply to these internally generated aspects of the tactical picture in much the same way as information items sourced from external information sources. In this case the certainty assigned to internally-generated information will need to take into account the certainty assigned to the information upon which the inference was based. Fortunately, the expert systems literature is replete with methods to deal with reasoning under conditions of uncertainty (e.g. see Stefik, 1995; Jackson, 1999)

3.5 Information Dissemination

In order to be useful the proposed system needs to recruit the support of outside agencies whose information processing objectives are commensurate with those of the system. The TDS will neither have the authority or wherewithal to implement humanitarian initiatives, rather its role is to facilitate the actions of executive agencies and knowledge workers

[23] These issues are particularly pertinent in the face of a recent large-scale humanitarian disaster resulting from an undersea earthquake in the Indian Ocean on 26th December 2004. The earthquake caused a tsunami that devastated coastal communities in Sri Lanka, Indonesia, Malaysia, India and Thailand resulting in the deaths of approximately 40000 individuals. While the earthquake was detected by analysts at the U.S. Geological Survey it was reported that it was not clear who should be contacted to forewarn of the impending disaster The tragedy prompted calls for an earthquake and tsunami monitoring system – as already deployed in the Pacific – to be set up in the Indian Ocean.

such as humanitarian aid operatives and military personnel. From the perspective of a larger problem-solving process, namely the planning, coordination and deployment of humanitarian relief operations, the system is expected to play a limited, albeit valuable, role in terms of improving the situational awareness upon which strategic decisions by executive agencies are made. The aims of the system are therefore only fulfilled via its inter-operation with other agents, which focuses attention on the importance of the interfaces with those agents (an issue that is more fully fleshed out in Section 3.6). There is an emerging consensus in the knowledge engineering community that knowledge systems can best achieve highly complex problem solving objectives by establishing cooperative alliances with information processing agents whose cognitive profiles are selectively suited to distinct aspects of the problem solving process. 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 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 time*. In order to fulfil this function the system will require background knowledge about the roles of external agents, their executive capabilities and constraints, responsibilities to act, position in power hierarchies and communicative profiles. 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.

The TDS can assist in this respect by facilitating the establishment of task-relevant communication networks, i.e. communication between agents that need to inter-operate in the context of some particular task. One example is based on the need to contact particular agents or agencies with respect to the occurrence of an humanitarian disaster event or helping an agent locate relevant agents based on the current task context and the target agents role characterization. Such communities of practice may either be explicit in the form of agent subscription or registration policies, or they may be largely implicit and derived from an agents task-specific interests and concerns. In addition, communities of practice may be characterized in terms of their flexibility (i.e. the extent to which a community of practice can be assembled on demand in light of ongoing tasks and events) and permanence (i.e. the extent to which a pattern of communication is stable across time). In all cases a community of practice will tend to crystallize around a common set of task commitments, which may only be fully realized by facilitating agent communication and enabling information transfer based improved situational awareness of events in the current operational environment.

Agent characterizations are also important for the TDS in terms of indicating the different information needs and requirements of individual agents. Each agent will have a particular set of epistemic requirements in order to fulfil its executive and decision-making responsibilities. These requirements determine the kind of information each agent wishes to receive, a fact the system can exploit to selectively filter the information provided to the agent. In this case we see a secondary attentional filter-like process that is specifically geared to the interests and concerns of each agent that subscribes to the system. Based on this information the TDS can proactively notify agents of events and information that may be of specific interest to them or it may elect to highlight certain aspects of the tactical picture so as to ensure that certain information items are not overlooked by the end user. Clearly, this latter capability relies on knowledge about the perceptual and cognitive biases of external agents, an issue which is addressed in more detail in Section 3.6. In general, the information dissemination capability of the system will assume the form of a monitoring function that continuously monitors the tactical picture and notifies external agents of events of potential relevance to their specific information processing objectives. Key to this ability is the provision of a semantically-enriched framework or common vocabulary within which epistemic needs and concerns can be expressed and communicated. The ontology engineering initiative to be undertaken in the context of the current project aims to deliver just such a framework. The framework will provide a set of conceptualizations which agents can use to express their interest in particular events. For instance, an organization may wish to be notified about the occurrence of events deemed to have a severe humanitarian hazard liability or, more specifically, of earthquake events deemed to have a severe humanitarian hazard liability.

3.6 Agent Interaction and Visualization

The manner in which information is presented to cognitive agents can have a profound impact on their subsequent information processing activities and decision-making capabilities. In its crudest form we can imagine the differential impact of a command line interface versus a rich graphical interface on the working practices and task outcomes of a human operator working at a computer terminal. The importance of the perceptual environment is often overlooked in cognitive analyses of problem-solving abilities, but given that every action is a response (either directly or indirectly) to a sensation, then, to some extent at least, it seems as though the solution to most problems must inhere in the information-bearing environmental structures that support such action. We can use the term ‘effective environment’ to describe the aspects of the perceptual world that propitiate the preferential selection of action patterns and modes of thinking that corresponding to the points along a response chain that ultimately leads to the successful resolution of some real-world problem. The notion is not new; it has its theoretical forbears in the psychology literature, particularly in areas such as attention where the appropriate selection of information inputs is essential to the expression of adaptive action patterns. If it was not for the ability to select and filter information along dimensions of relevance, animals would quickly risk sensory overload. Animals avoid sensory overload by selectively attending to those sources of perceptual input of greatest importance for their ongoing species-specific interests and concerns. In this way the frenetic pandemonium of the sensory world is reduced to just a handful of stimulus cues to which a behavioural response is justified; and in most cases the information processing resources of the organism are carefully constrained (by evolution) to respond in ways that facilitate the survival and long-term inclusive fitness of the individual. The notion of an ‘effective environment’ is similar to the concept of an animal’s ‘umwelt’ (see Clark, 1997; pg 24-25). An animal’s ‘umwelt’ is defined as the set of environmental parameters to which it is sensitised to respond to. It is those features of the environment which an organism is predisposed to process with high priority. Animals thus inhabit different 'effective environments' in which perception is skewed towards those features of the perceptual environment that matter to an organism in terms of its species-specific needs and concerns.

The key question raised by this discussion of 'effective environments' is how can we structure the information presented to cognitive agents in such a manner that it befits both their idiosyncratic information processing and perceptual profiles, while simultaneously supporting cognitive or motor responses that are intermediate steps on the path to a successful solution state for some domain-specific problem? This is a complex issue because the kinds of information-bearing environmental structures that are best suited to eliciting favourable response outcomes in one set of agents may not, and in general will not, be true of another set of agents. The issue is particularly problematic in the case of human operatives where a prolonged period of exposure or training may lead to a dependence on certain forms of perceptual input as a prerequisite for expert performance (consider the ways in which the subtle alteration of some feature of the environment may lead to a catastrophic collapse in expert performance[24]). It is likely that the executive agencies in the domain of humanitarian relief operations will each have their own rather idiosyncratic set of perceptual biases (some of which may have been acquired during extensive periods of training, e.g. military operatives). These biases and preferences determine the manner and, perhaps more importantly, the modality in which information should be presented. The system should therefore be capable of supporting multiple visualizations[25] of the tactical picture in a manner that is optimised for the perceptual and cognitive profiles of a variety of information consumers. In some cases the presentational format will be fixed as is the case with a military operator who has come to expect a standard visual layout and symbology (indeed such layouts may be sanctioned by military authorities). In other cases a more flexible strategy can be adopted in which the way information is presented to end users is entirely at the behest of the TDS. In both cases the system will require background knowledge about the presentational preferences and perceptual profiles of end-users and other inter-operating agents. This enables the system to support multiple, dynamic visualizations of the tactical picture, each of which is optimally suited to the processing capabilities of a particular type of end-user agent. In the extreme case one can imagine that information may be communicated using different modalities at different times depending on the current cognitive workload of an end user, e.g. the same information may be communicated using visual, aural or vibrotactile displays[26].

[24] An example from the psychology literature concerns the exploitation of persistent environmental cues by expert bartenders, in this case distinctively shaped glasses, to help recall and sequence orders in noisy and crowded environments. (Beach, 1988) Expert performance plummets in tests involving uniform glassware, whereas novice performances are unaffected by such subtle alterations of the perceptual working environment.

[25] The term 'visualization' is used here as a catch-all phrase to cover instances of information transfer which may not necessarily rely on the visual modality. One can imagine, for example, that certain types of information may be communicated via an auditory modality.

[26] This point touches on issues of future collaboration (discussed in Section 7). Past projects have addressed the selective presentation of information to combat pilots according to their current cognitive state (e.g. the Cognitive Cockpit programme). In addition, a number of other projects within the DIF DTC are currently focusing on the best way to reduce cognitive workload and operational efficiency by distributing information receipt across multiple modalities.

In addition to characterizations of agents in terms of their preferences or requirements for the receipt of information, the TDS will also need to pay close attention to the device through which information is to be communicated. Visual display devices, for example, vary according to their screen resolution, screen size, and colour support and these factors may influence the way in which information is to be presented. In order to address this concern the TDS will need access to semantically-detailed characterizations of the display capabilities of different devices used for communication and information exchange between a variety of cognitively diverse agents. Clearly, when one considers the delicate interplay between the characteristics of display devices, the cognitive profiles and perceptual biases of end-user agents as well as the need to support multiple 'effective environments' in which information can be used to support enhanced situational awareness and propitiate operationally useful decision outcomes, the role for complex knowledge processing and decision support becomes of paramount significance.

3.7 Summary

This section described the capabilities of the prospective system with respect to a number of capability areas. Such functional requirements serve to impose constraints regarding the specific set of tasks to be implemented by the prospective system and, as such, the current section serves as a precursor to task model specifications that detail the various tasks, functions and sub-functions to be implemented by the system. The need to provide a task-oriented decomposition of the prospective system should be undertaken in subsequent modelling initiatives and the recommendation is to rely on the Task Model component of the CommonKADS model suite for this particular purpose.

4 Technological Infrastructure

4.1 Overview

This section provides an overview of the range of technologies that will be exploited in some form in order to fulfil the capability requirements discussed in Section 3. The list of technologies reviewed herein is not intended to be exhaustive; other technologies may be exploited at later stages of the project.

4.2 Semantic Web Technologies

4.2.1 RDF / RDFS

RDF[27] is a general framework for describing a resource's metadata, or the information about the information available from the resource. The basic concept behind RDF is that a Resource is described through a collection of Properties called an RDF Description. Each of these Properties has a Property Type and Value. Any resource can be described with RDF as long as the resource is identifiable with a URI. An RDF model can be serialized to XML using a set of representational conventions specifically adopted for RDF (see Figure 4-1).

```
<? xml version="1.0" ?>
<RDF xmlns = "http://w3.org/TR/1999/PR-rdf-syntax-19990105#"
     xmlns:DC = "http://purl.org/DC#" >

  <Description about = "http://dstc.com.au/report.html" >
    <DC:Title> The Future of Metadata </DC:Title>
    <DC:Creator> Jacky Crystal </DC:Creator>
    <DC:Date> 1998-01-01 </DC:Date>
    <DC:Subject> Metadata, RDF, Dublin Core </DC:Subject>
  </Description>
</RDF>
```

Figure 4-1:RDF/XML Syntax

RDF relies on a particular vocabulary to characterize a set of resources. Such vocabularies can be defined using a schema language for RDF vocabularies: RDF(S)[28]. RDF(S) can be used to define the meaning, characteristics and relationships of a set of properties used in the characterization of a resource. This may include constraints on potential values and inheritance of properties from other schemas. The RDF schema is based on the same model as the RDF syntax specification and as such can be serialized as valid RDF/XML (see Figure 4-2)

27 http://www.w3.org/RDF/

28 http://www.w3.org/TR/rdf-schema/

```
<? xml version="1.0" ?>
<RDF  xmlns = "http://w3.org/TR/PR-rdf-syntax#"
      xmlns:RDFS = "http://w3.org/TR/WD-rdf-schema#" >

  <Description ID = "Title" >
    <type resource = "http://w3.org/TR/PR-rdf-syntax#Property" />
    <RDFS:label> Title  </RDFS:label>
    <RDFS:comment> The name given to the resource, usually by the
                   Creator or Publisher </RDFS:comment>
  </Description>

  <Description ID = "Creator" >
    <type resource = "http://w3.org/TR/PR-rdf-syntax#Property" />
    <RDFS:label> Author or Creator  </RDFS:label>
    <RDFS:comment> The person or organisation primarily responsible for
                   the intellectual content of the resource </RDFS:comment>
  </Description>

</RDF>
```

Figure 4-2:RDF(S) Syntax

4.2.2 RDQL

RDQL[29] is a query language for RDF models. RDQL can be used to query RDF repositories in a manner similar to the way in which SQL can be used to query conventional relational databases.

4.2.3 Jena

Jena[30] is an API and toolkit for building semantic web applications. It was developed by the HP Labs Semantic Web research group[31] and is open source; the source code can be downloaded from the HP website (http://www.hpl.hp.com/semweb/). Jena provides a programmatic environment for RDF, RDFS and OWL, including a rule-based inference engine, and is, in all likelihood the most popular API for the manipulation of RDF and OWL models. Jena allows a user to create and manipulate RDF models either in memory or within a relation database such as MySQL or Oracle. The latter mode of working uses the `ModelRDB` Java class to persist RDF data to a database and as such is similar to the mechanism used by 3Store technology (see Section 4.2.4).

Data within an RDF model can be accessed via Jena using specific API calls or by using RDQL. Jena provides support for RDQL via a query object (`Query`). Once instantiated, a query object can be passed to a query engine (`QueryEngine`) and the results stored in a query result object (`QueryResults`). Interestingly, this mechanism works irrespective of whether the RDF model is stored within a database or as a memory-resident object graph.

4.2.4 3Store

3Store was developed in the context of the AKT initiative at the University of Southampton (Harris & Gibbins 2003). It combines an RDF TripleStore with a query engine that promotes the efficient storage and retrieval of RDF metadata. The 3Store is implemented on top of MySQL database engine, which can be manipulated using

---

[29] http://www.w3.org/Submission/RDQL/

[30] http://jena.sourceforge.net/

[31] http://www.hpl.hp.com/semweb/

conventional queries formulated in SQL. However, in order to provide more sophisticated query capabilities, the 3Store incorporates an RDQL 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. The 3Store technology has been used successfully in a number of semantic web applications, including CS AKTiveSpace (see Section 4.8.2) and FloodSim (see Section 4.8.1).

4.2.5 OWL

OWL[32] is a W3C endorsed language for representing ontologies, where an ontology is defined as:

> "An ontology formally defines a common set of terms that are used to describe and represent a domain. Ontologies can be used by automated tools to power advanced services such as more accurate Web search, intelligent software agents and knowledge management." (OWL Web Ontology Language Use Cases and Requirements, 2004)[33]

The OWL development initiative was rooted in the DAML (DARPA Agent Markup Language) project – specifically the ontology language originating from this project: DAML+OIL[34]. OWL builds on RDFS to provide additional constraints that increase the semantic resolution of vocabulary descriptions. In particular OWL adds to the vocabulary used to describe properties and classes, including, relations between classes (e.g. disjointedness), cardinality, equality, richer typing of properties, characteristics of properties and enumerated properties.

OWL is available in three versions, each of which differs with respect to their formal complexity:

- **OWL Lite**: supports those users primarily needing a classification hierarchy and simple constraints. For example, while it supports cardinality constraints, it only permits cardinality values of 0 or 1. OWL Lite also has a lower formal complexity than OWL DL.
- **OWL DL**: supports those users who want maximum expressiveness while retaining computational completeness (all conclusions are guaranteed to be computable) and decidability (all computations will finish in finite time). OWL DL includes all OWL language constructs, but they can be used only under certain restrictions (for example, while a class may be a subclass of many classes, a class cannot be an instance of another class). OWL DL is so named due to its correspondence with description logics, a field of research that has studied the logics that constitute the formal foundation of OWL.
- **OWL Full**: supports maximum expressiveness and the syntactic freedom of RDF with no computational guarantees. For example, in OWL Full a class can be treated simultaneously as a collection of individuals and as an individual in its own right. OWL Full allows an ontology to augment the meaning of the pre-defined (RDF or OWL) vocabulary. It is unlikely that any reasoning software will be able to support complete reasoning for every feature of OWL Full.

---

[32] http://www.w3.org/2004/OWL/

[33] http://www.w3.org/TR/webont-req/

[34] see http://oiled.man.ac.uk/

OWL models can be represented as XML (see Figure 4-3), specifically as RDF/XML. This means that although OWL includes constructs that yield greater semantic expressivity, its XML representation remains valid RDF/XML and can be both parsed by RDF parsers (e.g. Jena) and stored in RDF TripleStore databases (e.g. 3Store).

```
<rdf:RDF xmlns="http://www.csd.abdn.ac.uk/research/AgentCities/WeatherAgent/weather-ont.daml"
    xmlns:daml="http://www.daml.org/2001/03/daml+oil#"
    xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
    xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#">

    <daml:Ontology rdf:ID="metarOnt">
        <rdfs:comment>An ontology for use in converting METAR and TAF reports to DAML</rdfs:comment>
        <daml:versionInfo>$Id: metar1.0.daml,v 1.12 2001/12/04 16:11:17 aelkiss Exp $</daml:versionInfo>
        <daml:versionInfo>$Id: metar1.1.daml,v 1.12 2002/01/23 16:11:17 ggrimnes Exp $</daml:versionInfo>
    </daml:Ontology>

    <daml:Property rdf:ID="hasWeatherEvent">
        <rdfs:comment>Weather reports contain information about weather events.</rdfs:comment>
        <rdfs:domain rdf:resource="#WeatherReport"/>
        <rdfs:range rdf:resource="#WeatherEvent"/>
    </daml:Property>

    <daml:Property rdf:ID="hasReportDate">
        <rdfs:comment>Every weather report must have a timestamp</rdfs:comment>
        <rdfs:subPropertyOf rdf:resource="#hasWeatherDate"/>
        <rdfs:domain rdf:resource="#WeatherReport"/>
        <rdfs:label>day of the month report was made</rdfs:label>
    </daml:Property>

</rdf:RDF>
```

Figure 4-3:OWL RDF/XML Syntax

4.3 Scenario Instantiation Tools

An effective demonstration of the capabilities of the TDS relies on the ability to define, load and execute scenarios. Ideally, the system should be able to cope with a number of scenarios depicting a different set if events in a robust and reliable manner. Although a specific scenario has been developed to showcase the capabilities of the TDS (see section 2.1), the scenario only exists as a paper-based specification of the major events and information objects that must be handled by the TDS. In order to define scenarios that can be used by the TDS for demonstration purposes it is imperative that machine-readable characterizations of scenario information can be made available. Clearly, the availability of software tools specifically designed to assist in the specification of a scenario will be invaluable in this respect. Additionally, we recommend using XML as the data format used to describe scenario instances.

4.4 Methods and Techniques

4.4.1 CommonKADS Methodology

Over the past fifteen years a comprehensive methodology has emerged that helps guide the process of acquiring, modelling and documenting knowledge content. The CommonKADS methodology (Schreiber et al, 2000) originates from the need to build industry-quality knowledge systems on a large scale, in a structured, controllable, and repeatable way.

One of the central insights of the CommonKADS approach is the need to develop a variety of models, each of which focuses on an a particular aspect of the problem at hand (see Figure 4-4). There are six models in the full methodology, although the extent to which these models require detailed explication varies from application to application.

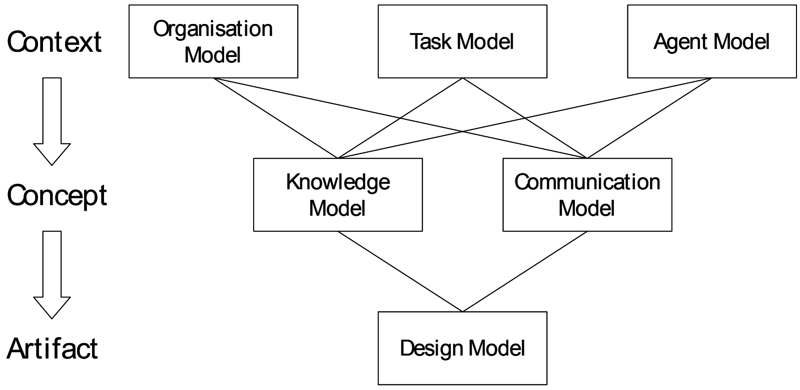

Figure 4-4:The CommonKADS Model Suite

Together the organisation, task, and agent models analyse the organisational environment and the corresponding critical success factors for a knowledge system. The knowledge and communications models yield the conceptual description of problem-solving functions and data that are to be handled and delivered by a knowledge system. The design model converts this into a technical specification that is the basis for software system implementation.

The CommonKADS methodology also exploits libraries of task templates which are derived from the notion of problem-solving methods (PSMs). For a particular type of task the task template specifies types of inference that are made in the context of the method, the control over inference execution and the knowledge roles that serve as the inputs and outputs of inferences (see Figure 4-5). The CommonKADS Methodology defines a library of task templates for a variety of different tasks including:

- **Analytic Tasks**
  - **Classification**
  - **Assessment**
  - **Diagnosis**
  - **Monitoring**
  - **Prediction**
- **Synthetic Tasks**
  - **Configuration Design**
  - **Modelling**
  - **Planning**
  - **Scheduling**
  - **Assignment**

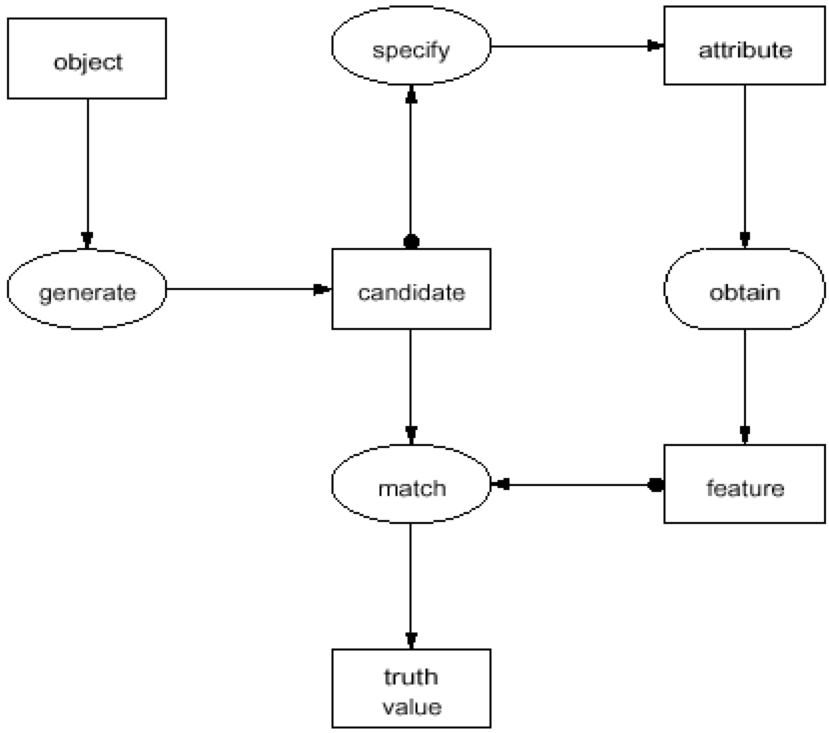

Figure 4-5:CommonKADS Classification Task Template

4.4.2 Knowledge Capture Techniques

The process of transforming domain-specific knowledge, much of which is implicit, into a set of symbolic and easily communicable representations is known as knowledge acquisition. Over the past two decades a number of specific techniques have been developed to aid in the elicitation and representation of knowledge. Such techniques range from so called 'natural' techniques such as interviews and behavioural analysis, through to more 'contrived' techniques such as concept sorting and repertory grids (Schreiber et al, 2000; Shadbolt et al, 1999).

Figure 4-6 illustrates a taxonomy of knowledge acquisition techniques. The intelligent deployment of such techniques can facilitate the process of eliciting and capturing knowledge, even if such knowledge is poorly documented or is not introspectively accessible, i.e. is implicit or tacit. Each of the above techniques is differentially suited for a particular type of knowledge acquisition context. The context, in this case, includes the type of domain expert to which the technique is applied (e.g. practitioner, academic), the form in which the knowledge is encoded (e.g. procedural, declarative), the extent to which the knowledge is explicit and therefore accessible via conscious introspection, and, finally, the type of knowledge which is to be elicited (examples knowledge types include concepts, attributes, rules, inference steps and the like). All these factors conspire to establish a strong 'differential access' hypothesis in which different knowledge acquisition techniques are differentially suited to the elicitation of particular types of knowledge in particular knowledge acquisition contexts.

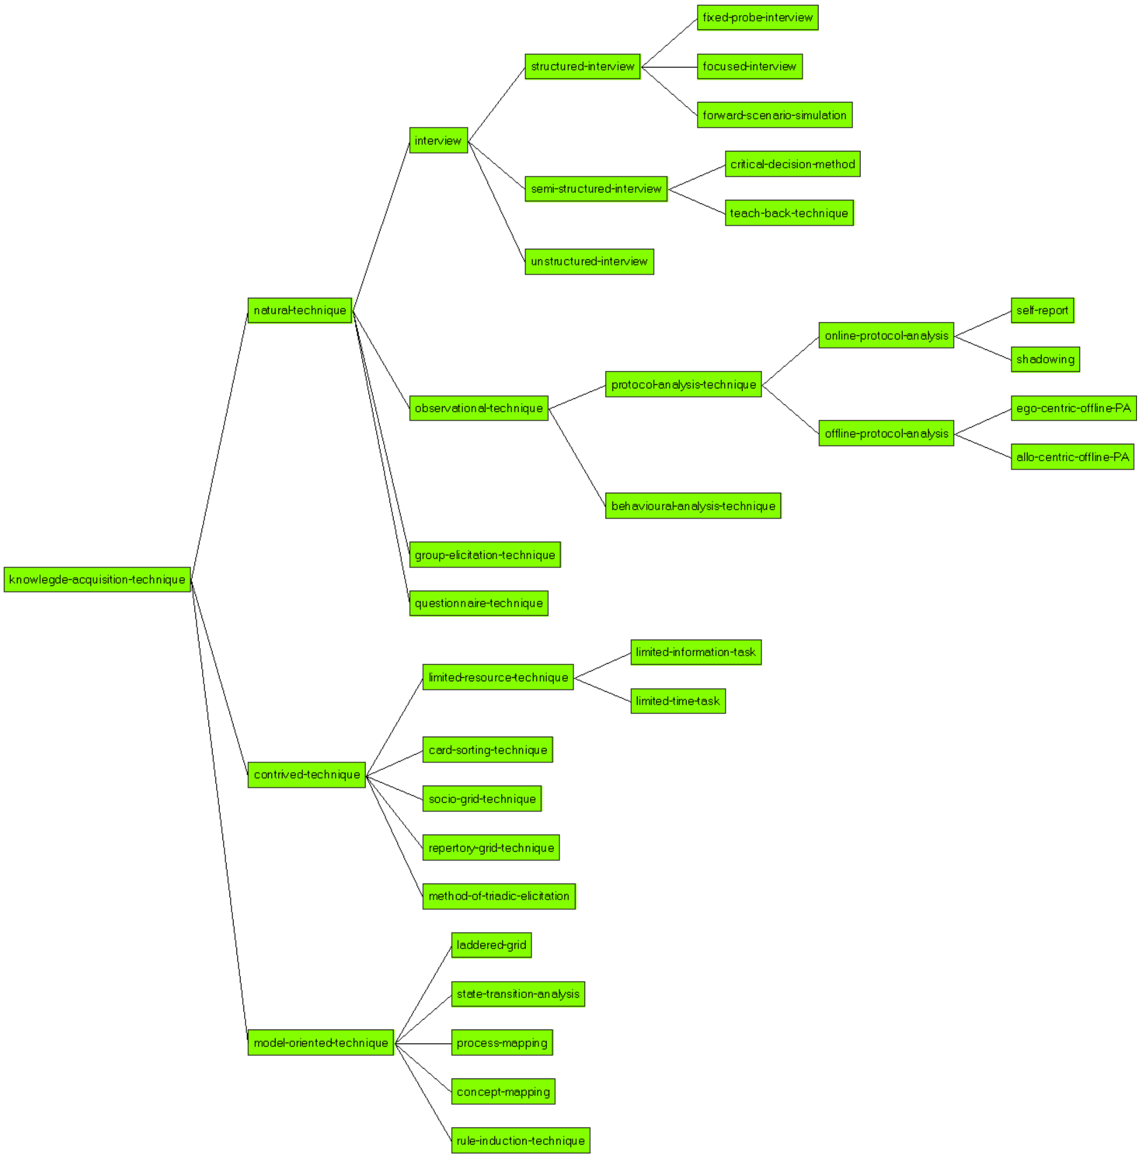

Figure 4-6: Taxonomy of Knowledge Acquisition Techniques

4.5 Knowledge and Ontology Editors

4.5.1 PCPACK / IKEW

PCPACK[35] and IKEW[36] represent specialized knowledge acquisition toolkits that may be used to capture knowledge from a variety of source materials such as textbooks, web pages, user manuals, interview transcripts and so on. They can also be used in conjunction with an SME to perform live knowledge acquisition in the course of a knowledge acquisition session. Both toolkits provide access to a range of tools that represent the software counterparts of a number of knowledge acquisition techniques (see

[35] http://www.epistemics.co.uk/Notes/55-0-0.htm

[36] http://www.i-kew.com

Section 4.4.2). Subsequent sections provide a brief overview of the tools available in these packages.

4.5.1.1 Laddered Grid Tool

The Laddered Grid Tool allows the knowledge engineer to construct laddered grids corresponding to decomposition hierarchies for a particular domain-specific relationship. The two most common types of hierarchies are taxonomic (is-a) and compositional (part-of); however, ladders can be constructed for any relationship (e.g. a functional decomposition), and mixed ladders (which feature multiple relation and element types) can also be used. The Laddered Grid Tool presents this hierarchical knowledge as a two-dimensional directional graph. The unique representational and visualization capabilities of the Laddered Grid Tool make it ideal for editing the hierarchical relationships between knowledge elements.

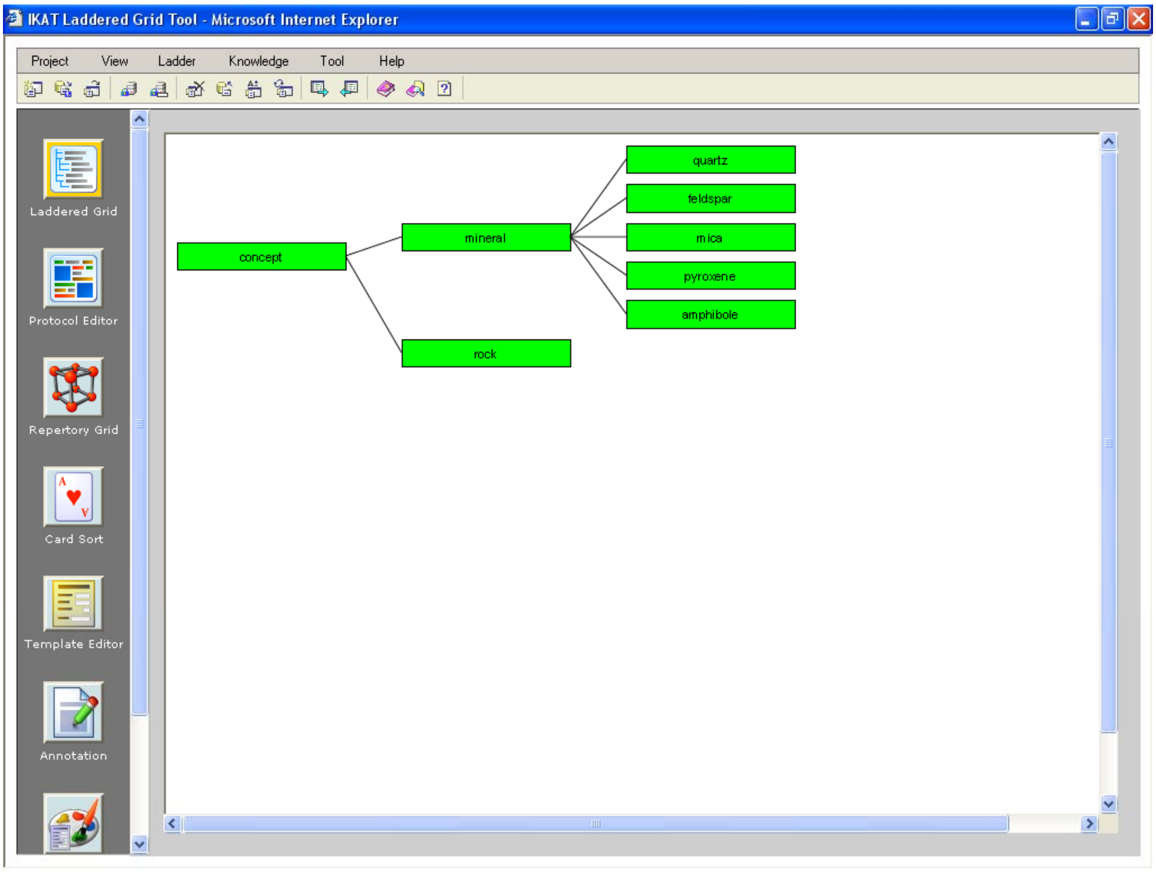

Figure 4-7:IKAT Laddered Grid Tool

4.5.1.2 Repertory Grid Tool

The Repertory Grid Tool is an implementation of a psychological technique used for revealing hidden or implicit conceptual structures. In its basic form, the expert is asked to discriminate between triads of entities in a domain, by suggesting a construct that applies to two of the entities but not the third. The resulting construct is then applied to other entities in the domain. Random triads can be presented, and new constructs can be created to discriminate between them. The result is a series of constructs that can be used to describe the domain. The contrived format of the repertory-grid technique means that some of the constructs might not even be apparent to the domain expert. A variety of cluster analytic techniques can be applied to the completed repertory grid as the basis for revealing the relative similarity of elements and constructs. The hierarchical clustering of elements can reveal complex taxonomies that may not have been consciously accessible to the expert.

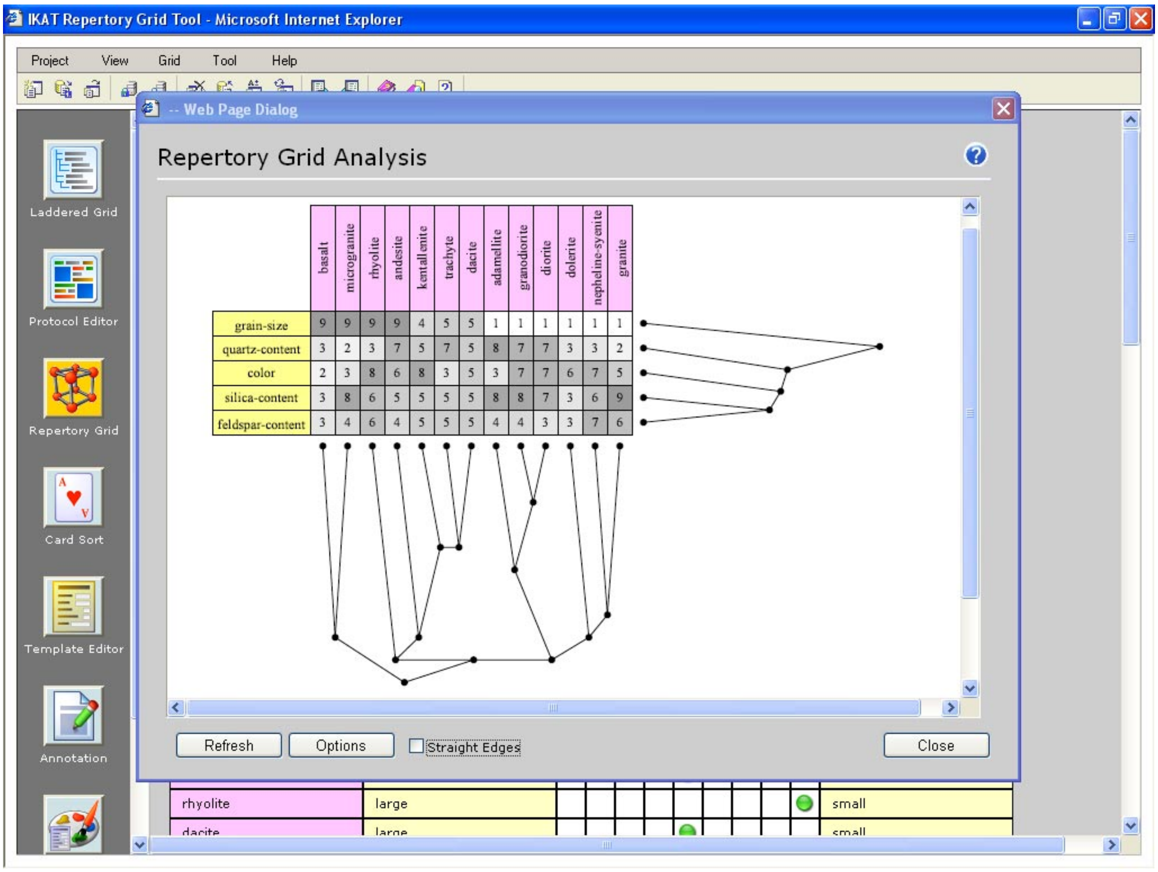

Figure 4-8:IKAT Repertory Grid Tool

4.5.1.3 Card Sorting Tool

The Card Sort Tool is a tool for sorting and categorizing groups of knowledge objects. The basic aim is to support the expert's grouping of concepts into significant clusters. The visual metaphor of the Card Sort Tool is that of taking a series of cards, each of which is associated with a visual representation of a knowledge object, and then sorting them into piles to make distinctions between them. By doing a card sort, the various ways in which experts 'see' a set of concepts, and the various types of distinction that they make between them, can be made clear.

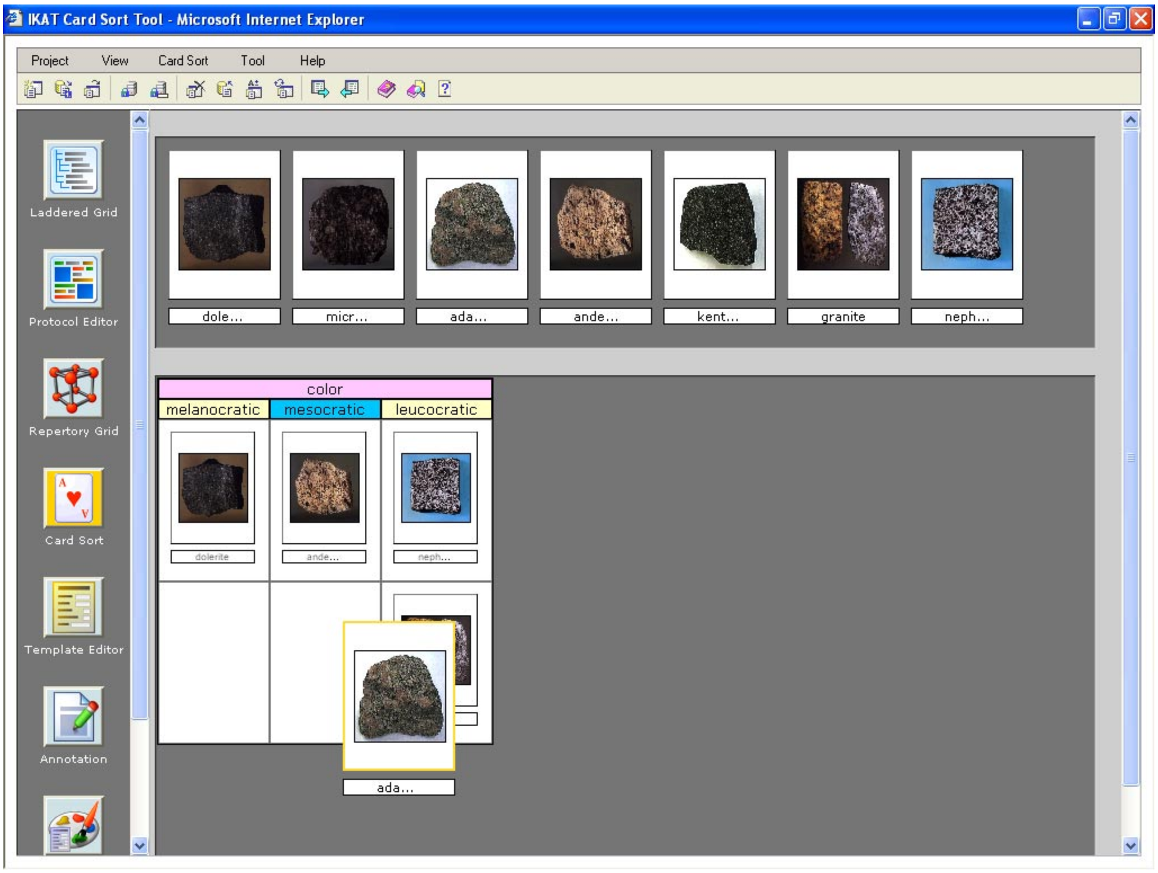

Figure 4-9:IKAT Card Sorting Tool

4.5.1.4 Protocol Editor Tool

The Protocol Editor is a tool used to annotate knowledge-rich textual resources, otherwise known as protocols. Protocols are important sources for knowledge acquisition; they can include books, documents, manuals, or transcriptions of experts' commentaries or interviews. Phrases in the protocol can be marked up in the Protocol Editor using special marker 'pens'. These allow the user to note occurrences of, among other things, concepts, attributes, values, relations, processes or rules, or to make notes connected with particular phrases. The various phrases isolated during the marking up period can be used to develop a lexicon, or glossary, of important terms in the domain. The knowledge objects that are isolated by the editing process can also be used at an early stage in a project as raw material for other knowledge acquisition tools.

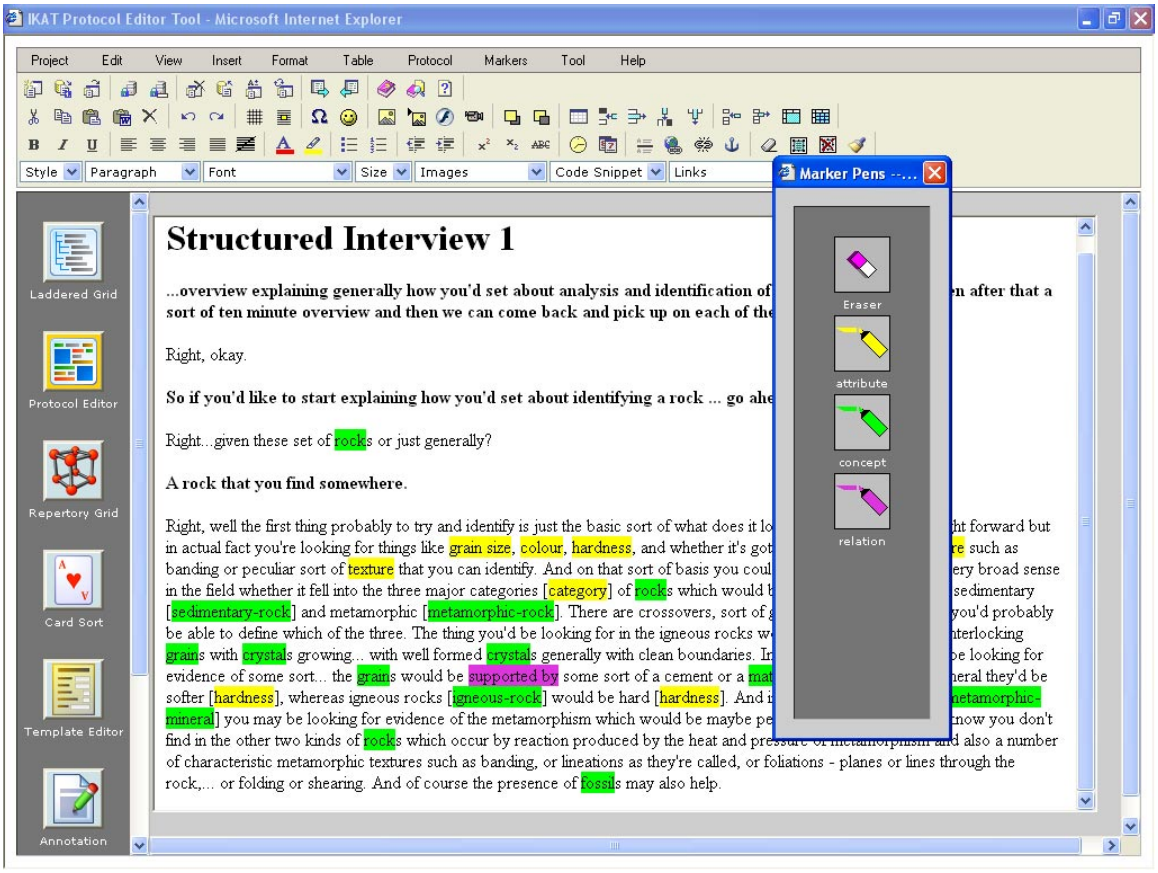

Figure 4-10:IKAT Protocol Editor Tool

4.5.1.5 Annotation Template Tool

The Annotation Template Tool allows a knowledge engineer to specify a template for a category of knowledge objects or a specific type of knowledge object, e.g. a specialization of a higher order concept. The template is designed using standard HTML technology and can thus support a potentially limitless variety of page layouts and presentational formats for information content. In addition, a number of application-specific formulas can be inserted into the template as the basis for automatically populating annotation pages with information about the knowledge object based on the knowledge infrastructure of the project, e.g. object characterizations and inter-relationships with other objects.

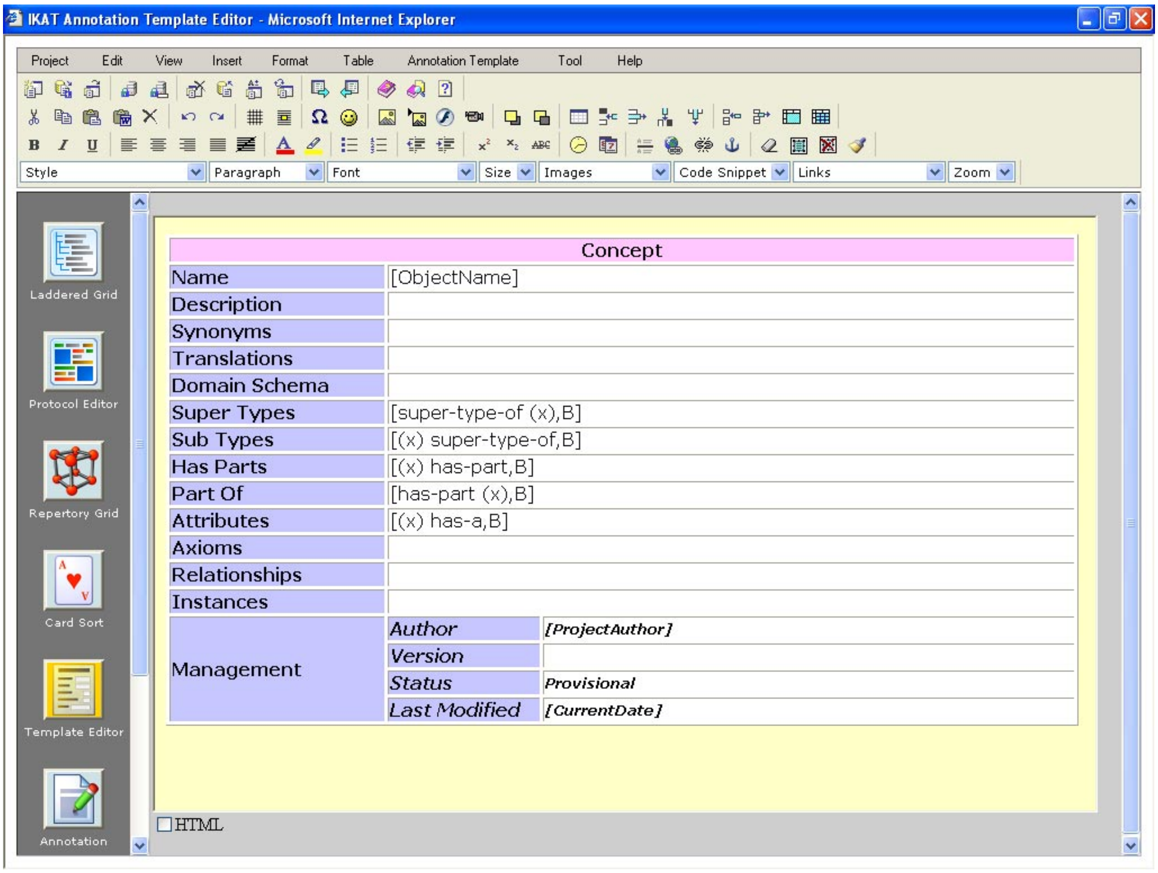

Figure 4-11:IKAT Annotation Template Tool

4.5.1.6 Annotation Tool

The Annotation Tool enables users to add annotations to their projects. An annotation page is essentially a standard HTML page that can be edited in a manner similar to the HTML editing capabilities provided by well-known graphical web authoring tools. Annotations can store explanatory text and multimedia information for any knowledge object which features in the knowledge base. Since the annotation page is an HTML page it is easy to navigate to related knowledge objects using a native hyperlinking capability. Each annotation page is derived from an annotation template (see Section 4.5.1.5), from which it inherits stylistic and layout information. The formulas defined in the underlying template are executed in the context of the annotation page in order to reflect the information infrastructure of a knowledge project. Such formulas are updated with every modification to the knowledge project, so the content of the annotation page always reflects the current project structure.

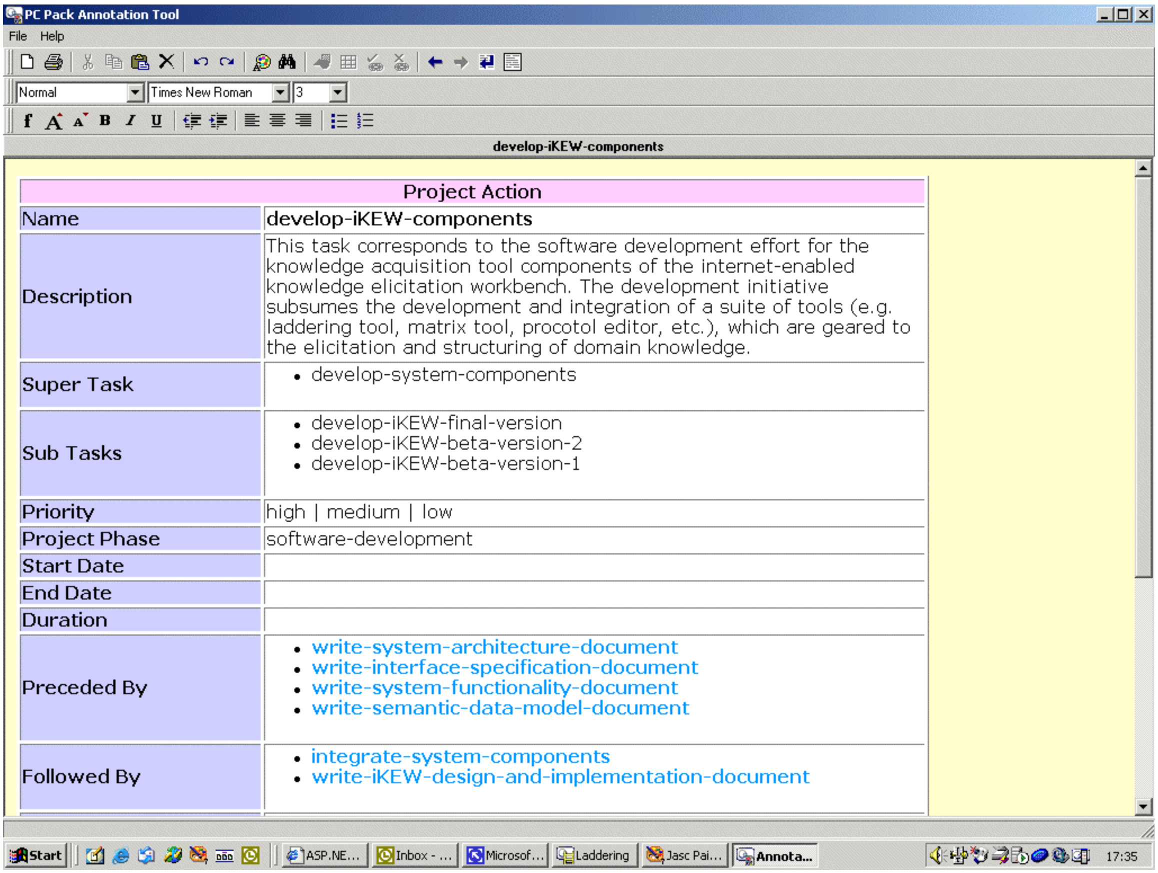

Figure 4-12:PCPACK Annotation Tool

4.5.1.7 Diagram Template Tool

The Diagram Template Tool can be used to specify a template for a particular type of diagram that is subsequently created and used in the context of the Diagramming Tool (see below). The Diagram Template Tool allows a user to place restrictions on the types of knowledge objects that can feature in a diagram as well as specifying the type of diagrammatic notation used to represent a knowledge object.

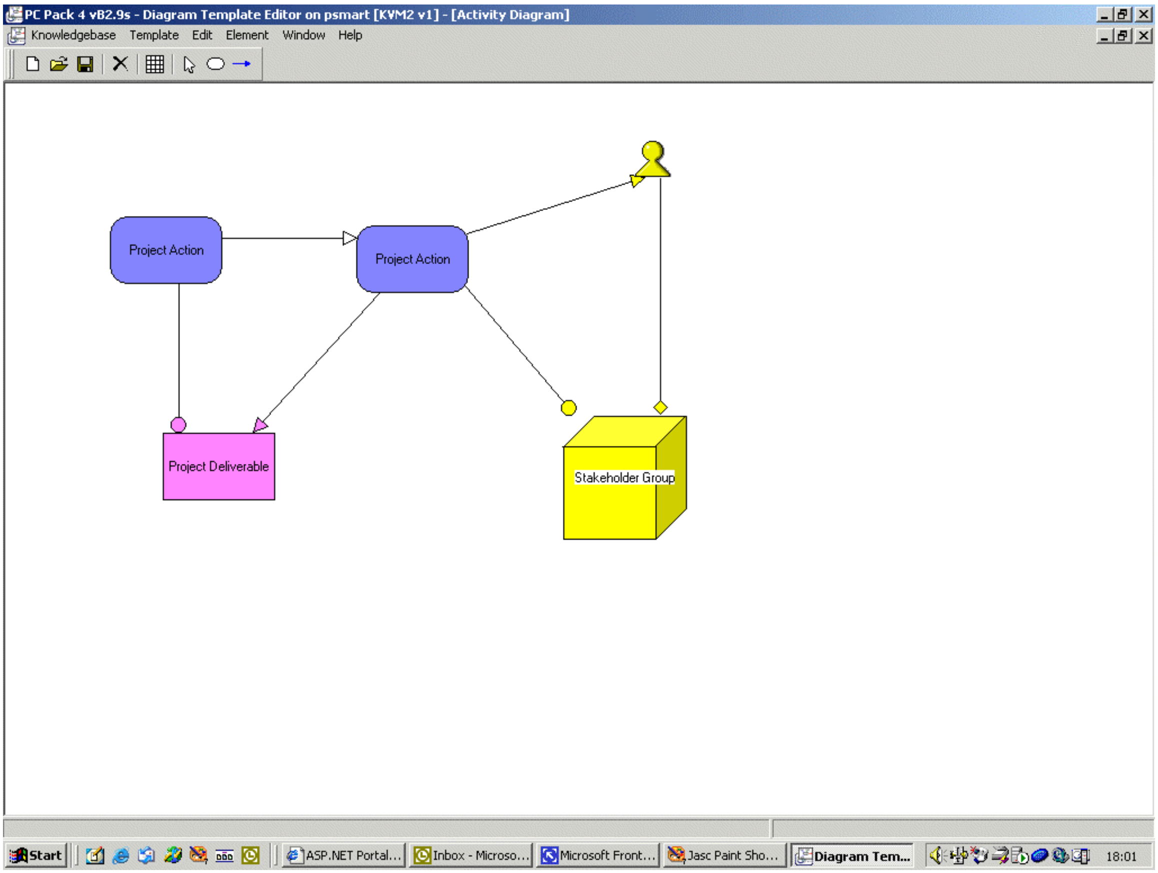

Figure 4-13:PCPACK Diagram Template Tool

4.5.1.8 Diagramming Tool

The Diagramming Tool allows the user to characterize and edit knowledge objects in a visually oriented manner using a predefined set of graphical modelling notations. Many modelling frameworks avail themselves of standard sets of graphical formalisms, which can be used to communicate modelling decisions in a concise and effective manner, e.g. the UML. The Diagramming Tool allows a user to specify any set of graphical notations, which characterize their preferred modelling methodology using the Diagram Template Editor (see above). Users can subsequently use these graphical modelling conventions to depict the inter-relationships between knowledge objects in a particular knowledge project.

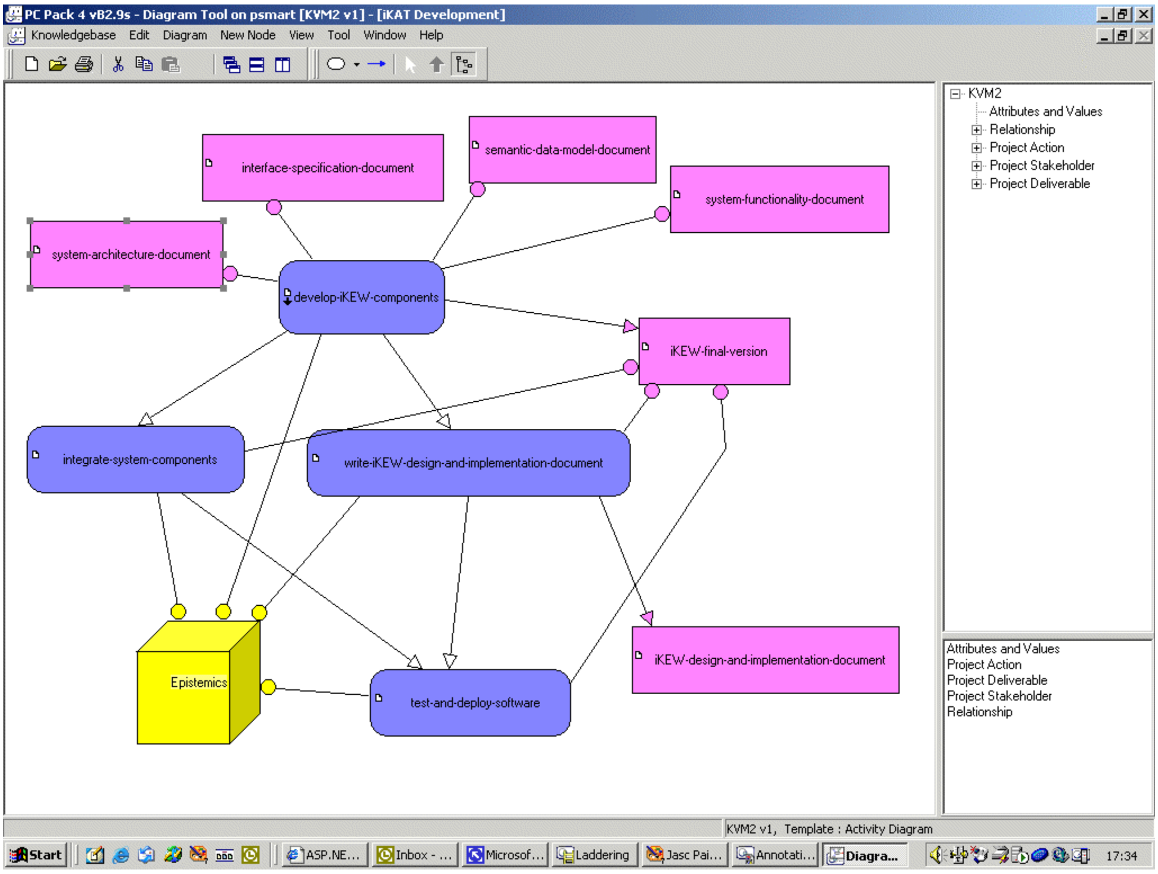

Figure 4-14:PCPACK Diagramming Tool

4.5.1.9 Matrix Tool

The Matrix Tool is a spreadsheet-like tool that can be used to assert relationships between knowledge objects, or to characterize knowledge objects in terms of attributes and values. The Matrix Tool provides an alternative and sometimes more intuitive editing environment for the elicitation and representation of human knowledge.

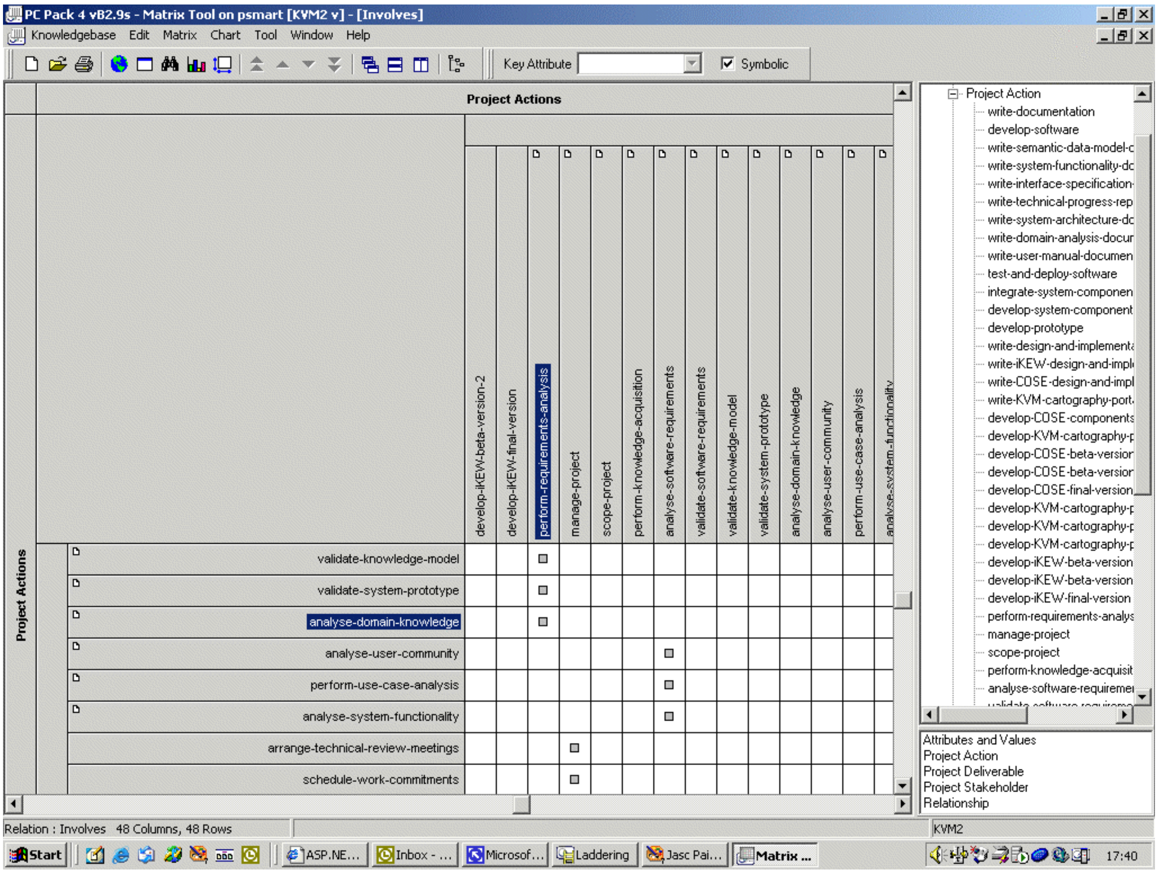

Figure 4-15:PCPACK Matrix Tool

4.5.2 eKADS

The eKADS CommonKADS Knowledge Editor[37] is a feature-rich knowledge editing environment specifically designed to support the CommonKADS[38] approach to knowledge engineering and management. It is primarily aimed at those in the knowledge engineering community who want to adopt a structured approach to the acquisition, representation and management of corporate knowledge assets. It provides a simple, easy-to-use forms-based environment for the specification of all aspects of the CommonKADS model suite. Once specified, model contents can be visualized as either CML (Conceptual Modelling Language) or XML (the native data format) (see Figure 4-16).

[37] http://www.epistemics.co.uk/ekads/

[38] http://www.commonkads.uva.nl/

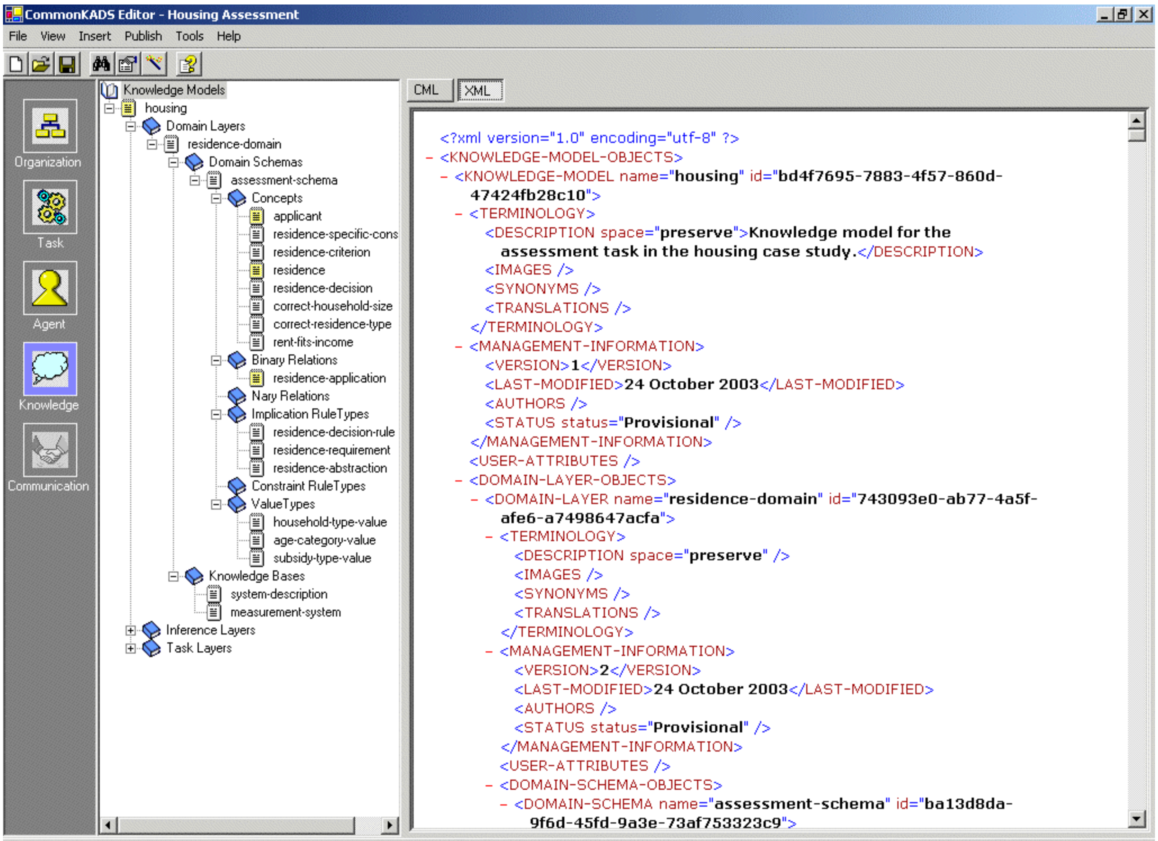

Figure 4-16:eKADS CommonKADS Knowledge Editor

4.5.3 Protégé

Protégé[39] is a Java-based ontology editor that provides a visual environment for the creation of ontologies and knowledge models. The Protégé interface (see Figure 4-17) provides access to a number of tabs which can be used to define class hierarchies and assert class properties. Protégé users can also design forms to facilitate the instantiation of ontology objects and perform queries against the data stored in the project. Completed models can be saved as plain text, knowledge webs, JDBC-accessible data stores and as RDF/XML.

[39] http://protege.stanford.edu/

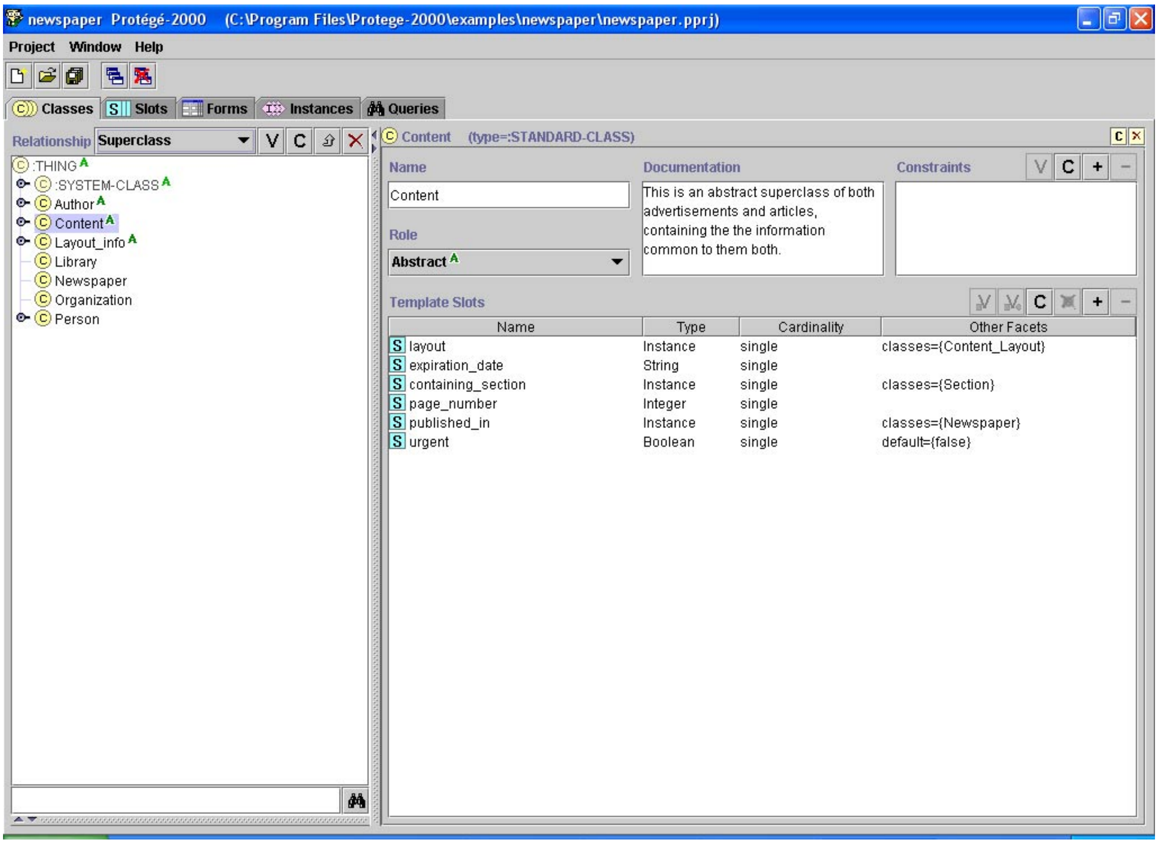

Figure 4-17:Protégé-2000 Interface

4.5.4 Isa-Viz

Isa-Viz[40] is a visual environment for browsing and authoring RDF models, represented as directed graphs (see Figure 4-18). Resources and literals are the nodes of the graph (ellipses and rectangles respectively), with properties represented as the edges linking these nodes. Since version 2.0, IsaViz supports GSS (Graph Stylesheets), a stylesheet language derived from CSS and SVG for styling models represented as node-link diagrams. IsaViz, as with most other RDF tools, is built on top of the Jena API (see Section 4.2.3.

[40] http://www.w3.org/2001/11/IsaViz/

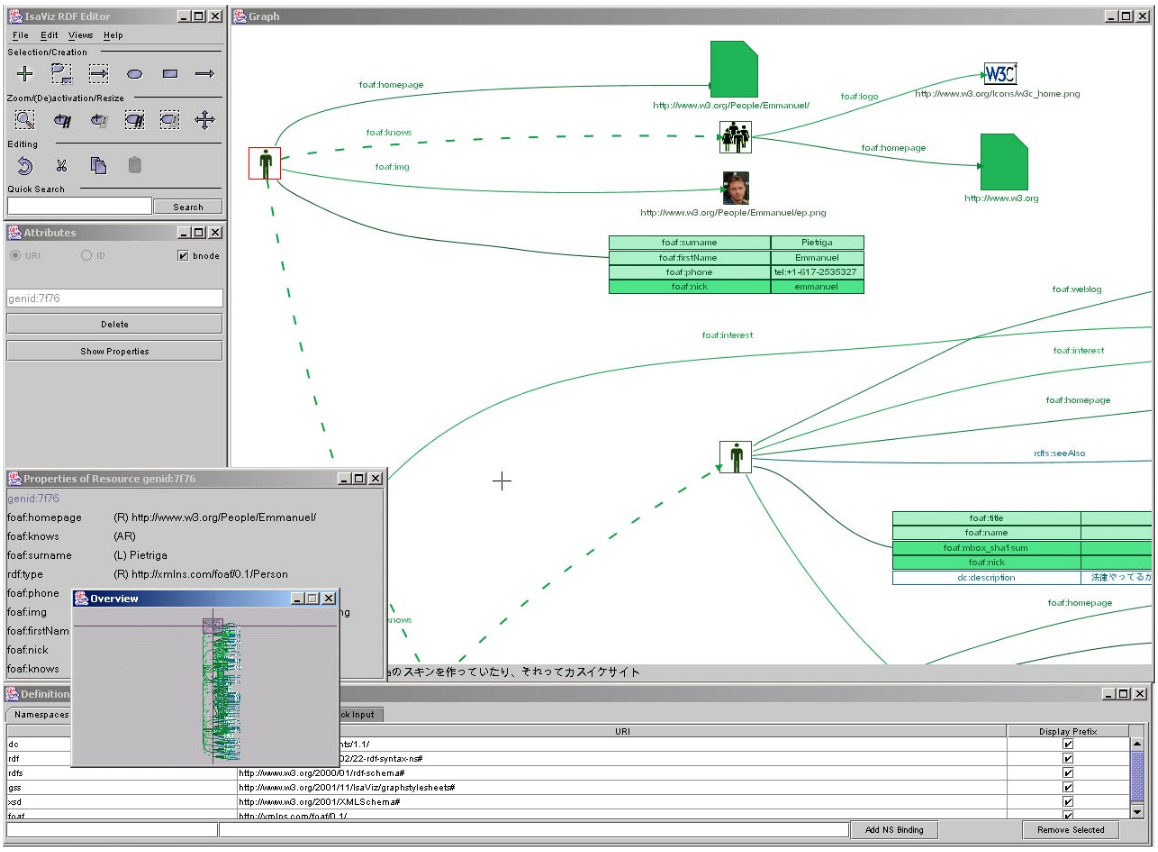

Figure 4-18:IsaViz Editor/Browser

4.6 Knowledge Representation Languages

4.6.1 OWL

See Section 4.2.5.

4.6.2 CLIPS / JESS

CLIPS[41] is an expert system shell that co-opts both a rule-based inference engine with object-oriented programming facilities. In terms of the current project CLIPS has a number of advantages over other expert-system shells:

- **cost**: CLIPS is free, which means that project resources can be spent on development time rather than software
- **size**: CLIPS is small, which means that it can be easily transported for testing and demonstration to other locations
- **system requirements**: CLIPS can be limited to run with little memory and few processor demands — while top-end computers give the best performance, it can run on low-end computers where necessary, which again gives benefits for transport to other locations for testing, integration and demonstration

[41] http://www.ghg.net/clips/CLIPS.html

- **interoperability**: CLIPS is multi-platform, which means that versions are available for different operating systems

The operation of CLIPS-based systems is characterised by the contingent activation of rules whose conditions are satisfied by the data structure of the reasoning environment. Usually, the data structure is specified by the properties of data objects which are instantiated from CLIPS classes, although rules can also operate over more simple CLIPS-specific data structures such as 'deffacts' and 'deftemplates'. In order for a knowledge system, implemented in CLIPS, to realize its reasoning objectives, the data objects within the CLIPS environment need to match the information structure of the environment in which the CLIPS system is embedded. This usually entails the creation of data objects within the CLIPS environment that faithfully match the characteristics of the external environment. Since we anticipate that, in the current project, the external environment will be represented by the contents of a central knowledge repository, viz. an RDF TripleStore, it is imperative that any decisions about the implementation strategy adopted for the current project should consider the inter-operation of the reasoning system with the knowledge repository components. In addressing this issue we propose the following alternatives:

1. **use of JESS (Java Expert System Shell) technology**: JESS[42] represents a Java implementation of CLIPS and incorporates JDBC technologies. Given the syntactic similarity of JESS to CLIPS, CLIPS code can be easily adapted for JESS based solutions. In addition, since JESS is open source, JESS can easily be extended and embedded within existing application environments.
2. **use of specialized code components**: since CLIPS is open source any CLIPS-based reasoning agent can easily be embedded in applications written using standard procedural languages, e.g. Delphi, Java, Visual Basic, C++, C#, etc. As such, one possibility is to embed CLIPS within a specialized application that handles all the interaction and communication capabilities of the CLIPS knowledge system with other software components.

In either case, the need for additional ancillary code components that specifically deal with inter-agent communication is a likely to be mandatory feature of the implementation strategy. Figure 4-19 illustrates a representative sample of CLIPS used to express an implication rule:

```
(defrule MAIN::combine-maximum-data "Combines data referring to the same instance-slot-value."
   (declare (salience 40))
   ?datum1 <- (object (is-a maximum-datum)
           (instance ?instance)
           (slot ?slot)
           (maximum ?value)
           (source $?sources1))
   ?datum2 <- (object (is-a maximum-datum)
           (instance ?instance)
           (slot ?slot)
           (maximum ?value)
           (source $?sources2))
   (test (neq ?datum1 ?datum2))
   =>
   (send ?datum2 delete)
   (send ?datum1 put-source $?sources1 $?sources2))
```

Figure 4-19:CLIPS Rule Syntax

42 http://herzberg.ca.sandia.gov/jess/

4.6.3 CommonKADS (CML)

The CommonKADS Methodology (see Section 4.4.1) avails itself of a custom modelling language designed to facilitate the communication of key features of the knowledge model developed for a particular problem domain. Figure 4-20 depicts a representative sample of this modelling language, also known as CML. The key advantage of CML is that it is a concise format for the communication of knowledge structures to a variety of different project stakeholders. The language is not overly concerned with the particular details of an implementation objective nor does it emphasize the logical detail epitomized by description logic languages such as OWL. Unfortunately, however, this is also the language's key weakness. Although CML parsers do exist[43], they are not in widespread use and are of little use in the absence of visual authoring tools that allow the automatic generation of CML. While this issue is being addressed by initiatives such as the eKADS initiative[44], there remains a notable absence of translators that can convert from CML to other representational formats, such as OWL.

---

[43] http://hcs.science.uva.nl/projects/void/intro.html

[44] http://www.epistemics.co.uk/ekads/

```
TASK classification;
  ROLES:
    INPUT: object: "Object that needs to be classified";
    OUTPUT: candidate-classes: "Classes consistent with the object";
END TASK classification;

TASK-METHOD prune-candidate-set;
  REALIZES: classification;
  DECOMPOSITION:
    INFERENCES: generate, specify, match;
    TRANSFER-FUNCTIONS: obtain;
  ROLES:
    INTERMEDIATE:
      class: "object class";
      attribute: "a descriptor for the object";
      new-feature: "a newly obtained attribute-value pair" ;
      current-feature-se: "the collection of features obtained";
      truth-value: "indicates whether the class is consistent with
        object features obtained during the reasoning process";
  CONTROL-STRUCTURE:
    WHILE NEW-SOLUTION generate(object -> class) DO
      candidate-classes := class ADD candidate-classes;
    END WHILE
    WHILE NEW-SOLUTION specify(candidate-classes -> attribute)
        AND SIZE candidate-classes > 1 DO
      obtain(attribute -> new-feature);
      current-feature-set := new-feature ADD current-feature-set;
      FOR-EACH class IN candidate-classes DO
        match(class + current-feature-set -> truth-value);
        IF truth-value == false
        THEN
          candidate-classes := candidate-classes SUBTRACT class;
        END IF
      END FOR-EACH
    END WHILE
END TASK-METHOD prune-candidate-set;
```

Figure 4-20:Task and Task Method CML Specification for Classification Task Template

4.6.4 SWRL

SWRL[45] (Semantic Web Rule Language) is a proposal for a rules-based language for the semantic web. It is the confluence of two key technologies: OWL and RuleML[46] (a sublanguage of the Rule Markup Language). The proposal aims to extend OWL axioms with rule-like structures that can be used for more complicated forms of inference than has heretofore been the case with OWL.

The proposed rules are of the form of an implication between an antecedent (body) and consequent (head). The intended meaning can be read as: whenever the conditions specified in the antecedent hold, then the conditions specified in the consequent must also hold. Both the antecedent (body) and consequent (head) consist of zero or more atoms.

---

[45] http://www.daml.org/2003/11/swrl/

[46] http://www.ruleml.org/

An empty antecedent is treated as trivially true (i.e. satisfied by every interpretation), so the consequent must also be satisfied by every interpretation; an empty consequent is treated as trivially false (i.e., not satisfied by any interpretation), so the antecedent must also not be satisfied by any interpretation. Multiple atoms are treated as a conjunction. An XML syntax is also given for these rules based on RuleML and the OWL XML syntax (see Figure 4-21).

```
<ruleml:imp>
  <ruleml:_body>
    <swrlx:individualPropertyAtom  swrlx:property="hasParent">
      <ruleml:var>x1</ruleml:var>
      <ruleml:var>x2</ruleml:var>
    </swrlx:individualPropertyAtom>
    <swrlx:individualPropertyAtom  swrlx:property="hasBrother">
      <ruleml:var>x2</ruleml:var>
      <ruleml:var>x3</ruleml:var>
    </swrlx:individualPropertyAtom>
  </ruleml:_body>
  <ruleml:_head>
    <swrlx:individualPropertyAtom  swrlx:property="hasUncle">
      <ruleml:var>x1</ruleml:var>
      <ruleml:var>x3</ruleml:var>
    </swrlx:individualPropertyAtom>
  </ruleml:_head>
</ruleml:imp>
```

Figure 4-21:SWRL XML Syntax

4.6.5 RDF / RDFS

See Section 4.2.1.

4.7 Knowledge Repositories and Data Stores

4.7.1 3Store

See Section 4.2.4.

4.7.2 MySQL

MySQL[47] is a free, open source database, commonly employed with most of the popular server-side scripting languages including PHP, JSP, and ASP. It is used as the database back-end of the 3Store technology described in Section 4.2.4.

[47] http://www.mysql.com/

4.8 Other Systems

4.8.1 FloodSim

FloodSim[48] (Gibbins et al, 2003) was developed as part of the AKT initiative[49] as a means of demonstrating the integration of ontologically-motivated DAML-S-based Web services with agent communication languages in the context of a humanitarian aid scenario. The timeline for the scenario includes a rapid flooding event which forces the creation of new relief camps, and a hostile event upon a relief convoy which requires military intervention and support. The system also contains a number of agents which generate reports on the state of the world (e.g. refugee movements, meteorological reports and forecasts) with differing degrees of certainty. The links between FloodSim and the current initiative should be clear and we hope to build on this initial functionality in the context of the current project.

4.8.2 CS AKTiveSpace

CS AKTiveSpace[50] (Shadbolt et al 2004a) is a semantic web explorer for investigating the Computer Science research domain in the United Kingdom. It combines information from multiple heterogeneous sources, such as published RDF sources, personal web pages, and data bases in order to provide an integrated view of this multidimensional space. The content is gathered on a continuous basis using a variety of methods including harvesting and scraping of publicly available data from institutional web sites (Leonard and Glaser, 2001), bulk translation from existing databases, and direct submissions by partner organizations, as well as other models for content acquisition. The content assumes the form of an OWL ontology, which relies on 3Store technology (see Section 4.2.4) for content storage and retrieval.

4.8.3 FOAEW Decision Support System

The decision support system for FOAEW provides reasoning support for AEW operators engaged in the task of planning a series of mission assignments to CAP aircraft in the context of an AEW sortie. The system was developed by Epistemics on behalf of QinetiQ Ltd in the context of the FOAEW initiative. The system exists as a set of knowledge structures describing the problem domain and the reasoning capabilities of the system. Scenarios are defined that include instances of the knowledge structures defined in the core library and both sets of knowledge structures are loaded into the knowledge system (implemented using the CLIPS expert shell) during system initialization. The knowledge system was used to generate a visualization of scenario events and decision output at each step in the scenario timeline using standard web technologies, i.e. HTML and JavaScript. The interface of the resulting demonstrator is presented in Figure 4-22.

---

[48] http://www.aktors.org/technologies/floodsim/

[49] http://www.aktors.org/

[50] http://triplestore.aktors.org/

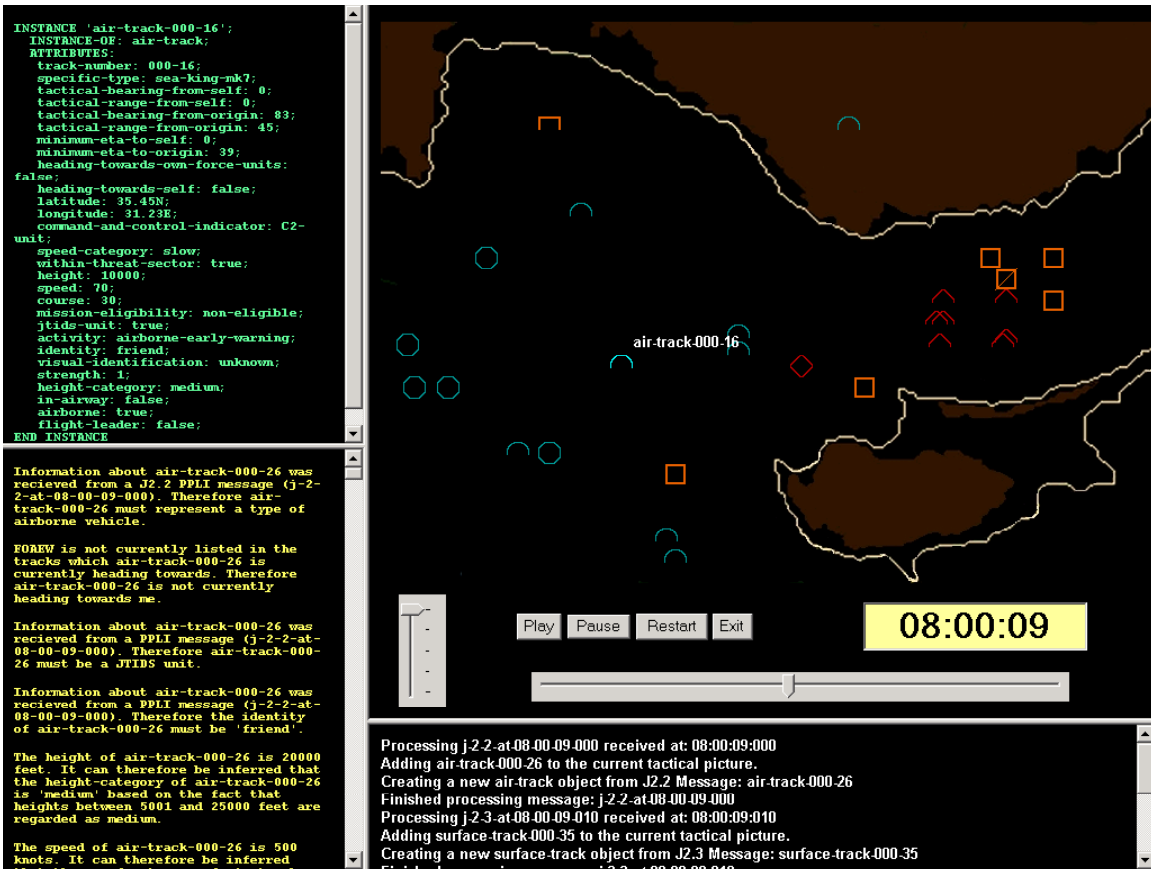

Figure 4-22:FOAEW Demonstrator Interface

The web interface features 4 panes (see Figure 4-22). The main screen (see Figure 4-23) depicts the actual track objects and geographical area in which AEW operations are occurring. The symbology of the track objects is consistent with that used by the Sea King Mk7 Mission System and the overall display characteristics of the main screen are intended to mimic the PPI on which the tactical picture is displayed. As the scenario progresses, the track objects move around the screen consistent with their changing latitude and longitude coordinates.

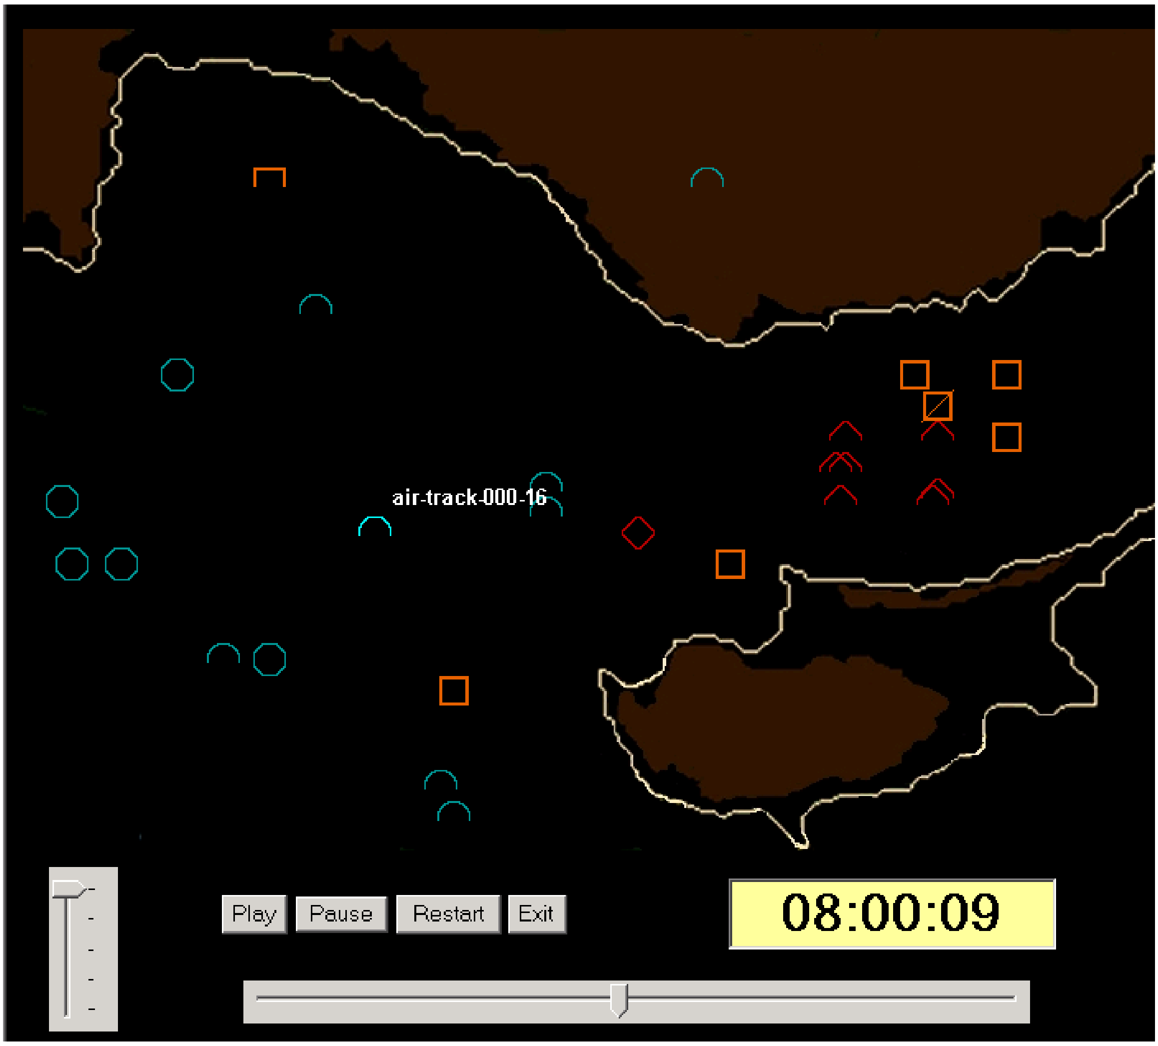

Figure 4-23: Main Display Window

The screen immediately below the main screen (see Figure 4-24) is used to provide ancillary information about the input/outputs of the knowledge system. In general, this screen will display the list of messages, which are received and processed by the knowledge system application at each point in the scenario timeline.

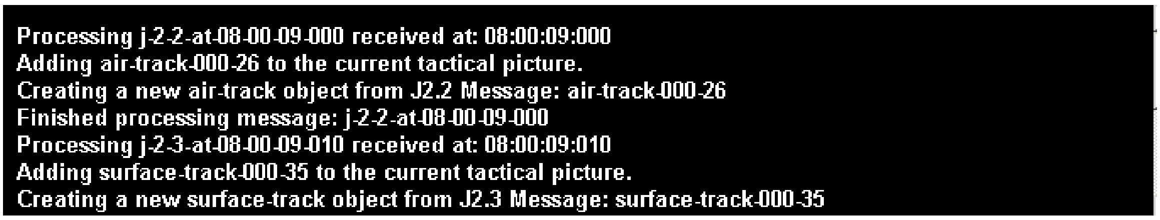
Processing j-2-2-at-08-00-09-000 received at: 08:00:09:000
Adding air-track-000-26 to the current tactical picture.
Creating a new air-track object from J2.2 Message: air-track-000-26
Finished processing message: j-2-2-at-08-00-09-000
Processing j-2-3-at-08-00-09-010 received at: 08:00:09:010
Adding surface-track-000-35 to the current tactical picture.
Creating a new surface-track object from J2.3 Message: surface-track-000-35

Figure 4-24: Message Processing Window

The screen to the top right of Figure 4-22, is the track information window (see Figure 4-25). Here, information about selected track objects is displayed using the standard notational format for instances of domain concepts. A track is selected by clicking on it with the mouse (the selected track is highlighted). Once a track is selected its associated information will continue to be displayed in the track information window until another track object is selected. Note that the information in this window is updated throughout the time course of the scenario.

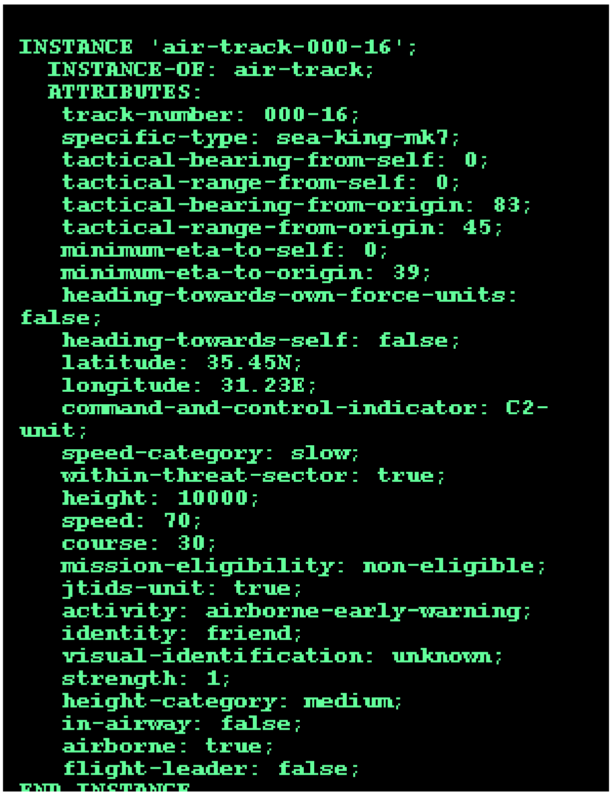

Figure 4-25: Track Information Window

The most important window, from the perspective of analysing and validating knowledge system output is found at the bottom left of the screen (see Figure 4-26). This window displays information about knowledge system output. This output can take the form of decisions, logical inferences or simply the provision of information about current processing status. The speed of the scenario can be slowed or paused entirely in order to more closely inspect the knowledge system output.

```
Information about air-track-000-26 was
recieved from a J2.2 PPLI message (j-2-
2-at-08-00-09-000). Therefore air-
track-000-26 must represent a type of
airborne vehicle.

FOAEW is not currently listed in the
tracks which air-track-000-26 is
currently heading towards. Therefore
air-track-000-26 is not currently
heading towards me.

Information about air-track-000-26 was
recieved from a PPLI message (j-2-2-at-
08-00-09-000). Therefore air-track-000-
26 must be a JTIDS unit.

Information about air-track-000-26 was
recieved from a PPLI message (j-2-2-at-
08-00-09-000). Therefore the identity
of air-track-000-26 must be 'friend'.

The height of air-track-000-26 is 20000
feet. It can therefore be inferred that
the height-category of air-track-000-26
is 'medium' based on the fact that
heights between 5001 and 25000 feet are
regarded as medium.

The speed of air-track-000-26 is 500
knots. It can therefore be inferred
```

Figure 4-26: KBS Output Window

As part of the need to instantiate diverse scenarios in a rapid and accurate manner a simple scenario generation utility was produced for the FOAEW project using Microsoft Excel in conjunction with VBA automation. The tool features the ability to script an entire scenario in terms of Link 16 Datalink messages. The information content of the messages can be specified using native Excel spreadsheet editing techniques and then exported as CLIPS data objects that represent the entire sequence of J series messages to be processed by the CLIPS application. The choice of Excel as an application development environment was inspired by the possible integration of these utilities with MANDRIL[51] (Stasys Ltd[52]). MANDRIL represents an integrated suite of software tools, the heart of which is a customised Excel application, with C++ extensions. It has the ability to interpret, decode and analyse J series message feeds obtained from live real-world recordings. MANDRIL can decode the information content of J series message feeds stored in a variety of data formats and thus facilitates the ability to convert actual message feeds into a format the CLIPS application could potentially process. Given that both MANDRIL and the scenario generation tools are produced in Excel, the task of integrating these utilities is relatively straightforward.

4.8.4 Decision Desktop

Decision Desktop (QinetiQ Ltd.) is a tool for visualizing information in a highly customisable fashion. The aim is to support multiple visualizations of information from different sources as a means to improve the decision-making competency of key

[51] http://www.mandril.co.uk/

[52] http://www.stasys.co.uk/

knowledge workers. The tool boasts a pluggable component architecture in which new views can be plugged-in as required without modifying the existing software.

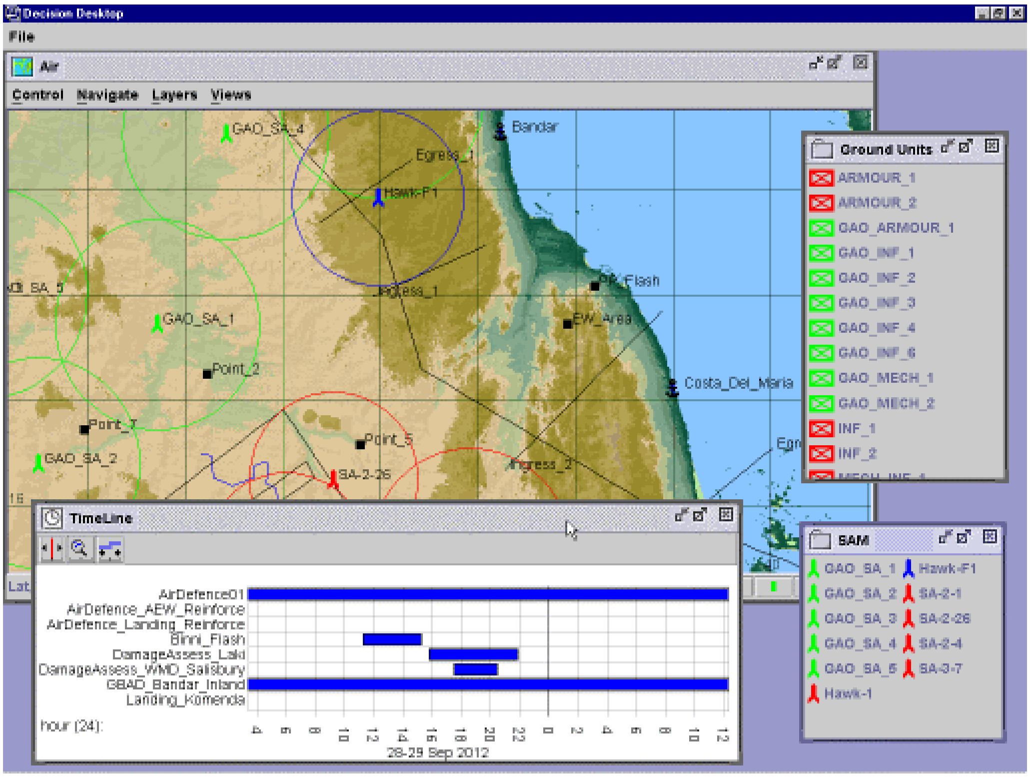

Figure 4-27:Decision Desktop Interface

5 Knowledge Infrastructure

5.1 Overview

This section provides an overview of the knowledge infrastructure of the target problem domain. The notion of knowledge infrastructure subsumes all knowledge structures, e.g. domain conceptualizations, problem-solving methods, knowledge-rich contingencies, etc., that play a role in ensuring the successful 'cognitive negotiation' of the problem domain. The term 'cognitive negotiation' is intended to reflect the ability of a problem-solving agent to realize inferential process, predicated on background domain-specific knowledge, that result in successful decision outcomes or the resolution of some problem-solving objective. The knowledge infrastructure essentially targets all those constructs of perceived relevance to the decision-making and reasoning capabilities of a prospective system. It typically provides a detailed, albeit partial, model of the kinds of reasoning processes and conceptualizations employed by human experts while thinking about or solving problems within a particular domain.

5.2 Knowledge Structures

Prior experience in knowledge engineering has revealed that a successful ingredient of knowledge infrastructure specification is the development of a number of knowledge-level (see Newell, 1980) models each of which targets a specific subset of the problem-solving constructs used in the domain. Contemporary approaches to knowledge engineering and management, such as CommonKADS (Schreiber et al, 2000), typically avail themselves of elaborate knowledge typing schemes, which are intended to make sensible contact with conventional software engineering methodologies, such as object-oriented programming, without reneging on the commitment to deliver high-fidelity, cognitively-oriented models of human expertise. The knowledge types typically encountered in the course of a knowledge modelling initiative include the following:

- **concepts**: A concept describes a set of objects or instances which occur in the application domain and which share similar characteristics. The notion of a concept is similar to what is called a 'class' in object-oriented modelling initiatives.
- **attributes**: Attributes represent the features or properties of a concept or relation., such as its size, colour, status, etc. At the level of instances, attributes can store atomic information, consistent with a particular type specification, that characterizes individual instances of concepts.
- **relations**: Relations describe the associations between concepts. The relation construct is typically employed when a characterization of one concept can only be specified with respect to a non-atomic data type, e.g. an instance of another concept.
- **tasks**: Tasks define the goals of an application. Typically task goals are specified in terms of input-output transitions.

- **inferences**: An inference is regarded as a primitive reasoning step which represents the lowest level of functional decomposition within the task model for a particular application. Typically, an inference uses knowledge contained in some knowledge base to derive new information from its dynamic input.
- **knowledge roles**: Inferences are indirectly related to domain knowledge constructs through the notion of knowledge roles. A knowledge role represents an abstract name for a collection of data objects that play a particular role in the reasoning processes of a system.
- **knowledge bases**: A knowledge base contains instances of the types specified in a domain schema, i.e. rules, instances and tuples. The knowledge base contains the actual data which is used by the reasoning system to realise the reasoning objectives of a knowledge-intensive task.
- **domain schemas**: A domain schema is a schematic description of the knowledge and information types relevant to particular application domain. A domain schema includes specifications of the concepts, relations and rule types that we find in a domain.
- **rule types**: The rule-type construct provides a means of describing sets of rules that share some common properties. There are no hard and fast rules about the criteria which should be used to group similar rule sets. Often rules which share the same role in a particular reasoning process should be grouped together. For example, rules of a particular type may all be relevant to a particular inference step within the reasoning process of the prospective system. The value of this heuristic is that it facilitates the process of updating the knowledge bases. Rules which share the same role in the reasoning process can be easily identified and edited within the knowledge base.
- **rules**: Rules define the constraints or knowledge-rich contingencies that exist between the attributes of concepts.

A knowledge model uses these constructs to represent distinct aspects of the target knowledge infrastructure using a variety of representational formalisms, e.g. natural language, formal modelling language, graphical notations, etc. The type of formalism selected for a particular task typically depends on a combination of the task objectives (e.g. system implementation vs. expert validation) and the type of stakeholder who participates in the task (e.g. system developer vs. domain expert). This is important since modern knowledge engineering initiatives typically involve teams of individuals each with different roles and different types of expertise. Ideally, knowledge models should flexibly support a complete range of representational and visualization requirements. In practice, however, the modelling commitments endorsed by different approaches (e.g. MOKA, CommonKADS, OWL) often complicates the full realization of this objective.

5.3 Knowledge Engineering Cycle

The current project has employed a tripartite approach to knowledge capture and knowledge model/ontology specification. Our approach consists of three distinct stages, including:

1. Knowledge Capture (Stage 1)

2. Knowledge Modelling (Stage 2)
3. Ontology Formalization (Stage 3)

Each stage of the process delivers outcomes that are utilized and refined in subsequent stages. In general, the knowledge structures encountered in each stage become progressively more refined and detailed as one moves from Stage 1 to Stage 3 (see Figure 5-1). However, it is important to bear in mind that the transition between stages is not necessarily assume a linear one, i.e. Knowledge Capture ⇒ Knowledge Modelling ⇒ Ontology Formalization; rather progress within one stage typically occurs in parallel with other stages. Iterative loops are established between the stages such that lessons learned from one stage feedback to influence the activities undertaken in the context of other stages (see Figure 5-2).

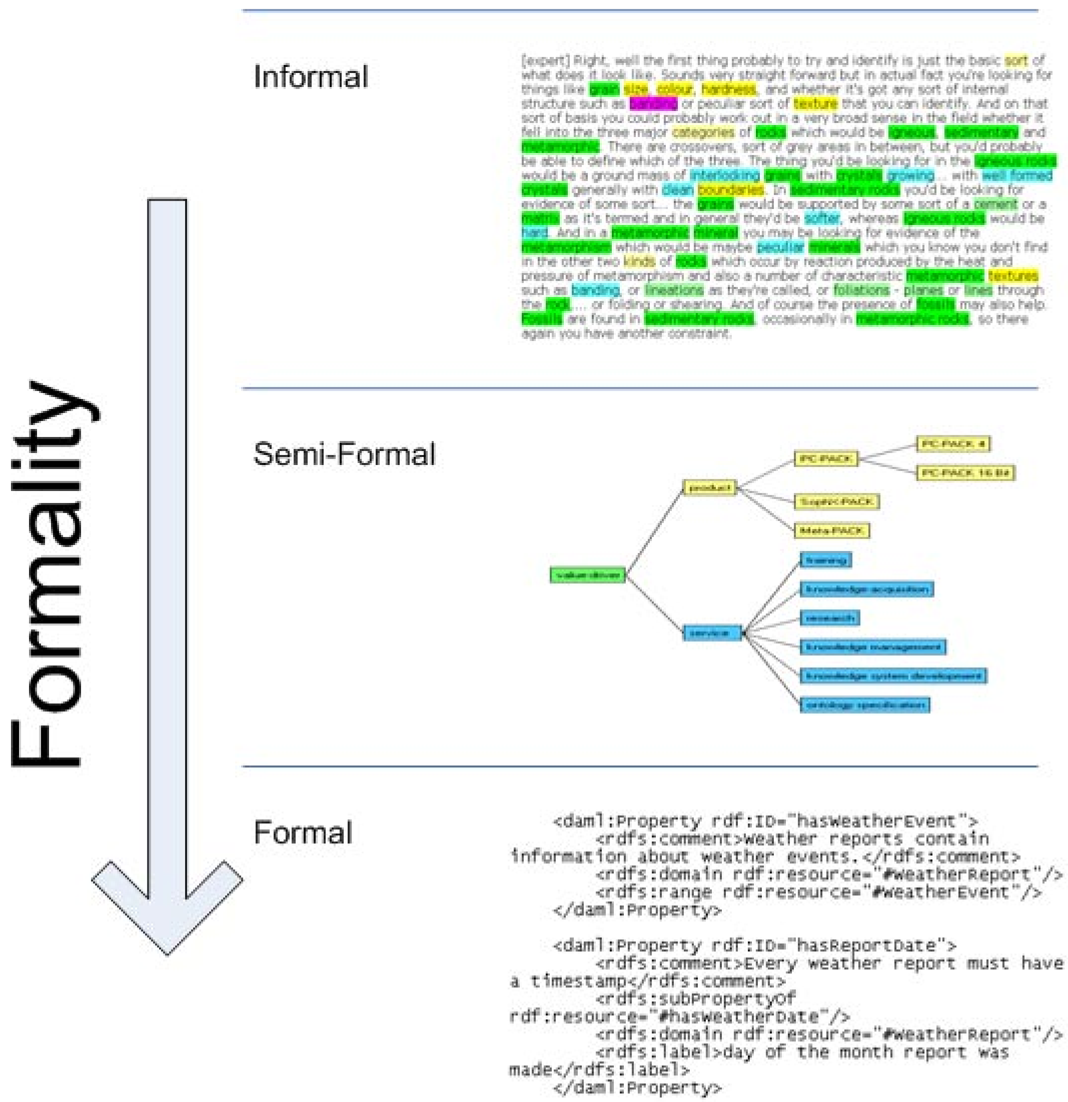

Figure 5-1:Levels of Formality

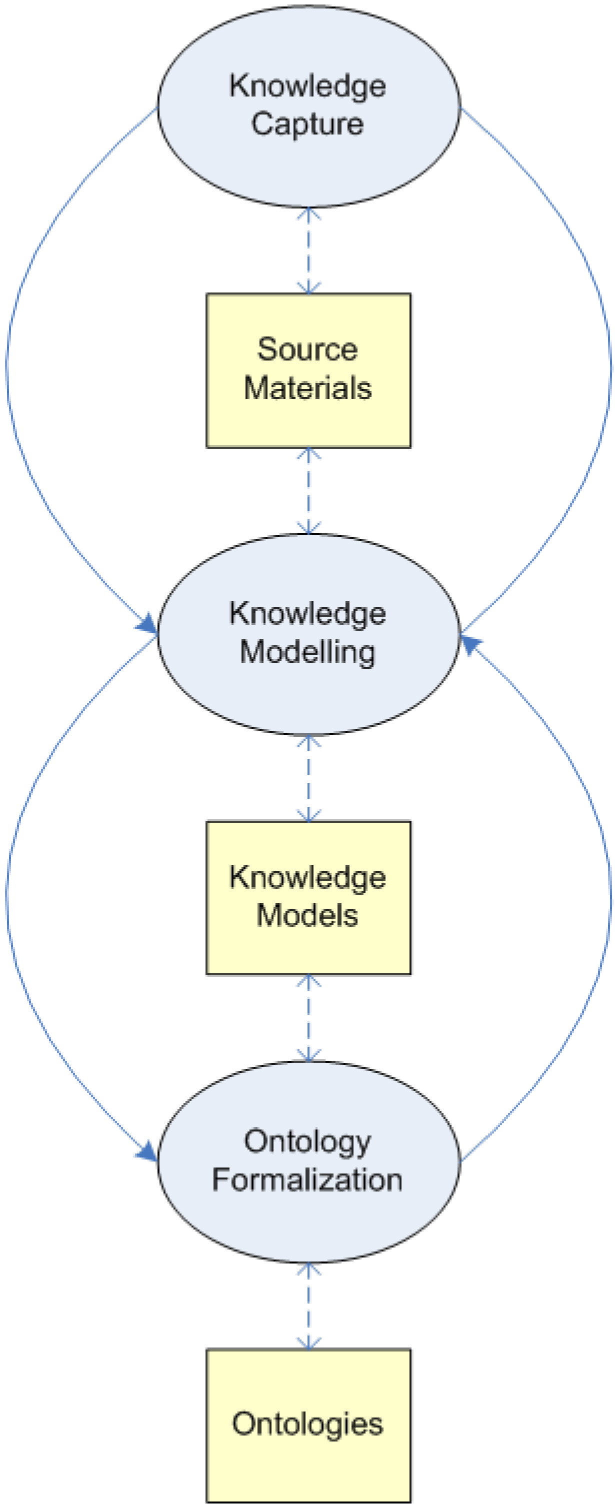

Figure 5-2:Knowledge Engineering Cycle

5.3.1 Knowledge Capture

In the first stage a variety of knowledge capture techniques (see Schreiber et al, 2000; Shadbolt et al, 1999) , e.g. protocol analysis, laddered grids, process modelling are used to acquire and elicit knowledge from a range of source materials, mostly public domain resources such as websites. In general the relevance of these resources to the current project was assessed based on whether the information content focused on one or more of the following areas:

- Humanitarian Operations/Agencies/Events
- Afghanistan (particularly sites with an emphasis on humanitarian or military affairs)
- Military Technology/Communications/Logistics

Based on this assessment we were able to identify a comprehensive list of source materials that served as the basis for subsequent knowledge capture initiatives. The outcome of this particular stage of knowledge modelling consists in a set of informal and semi-formal representations of domain knowledge.

5.3.2 Knowledge Modelling

The second stage of the knowledge engineering process involves the transformation of the knowledge structures acquired in Stage 1 to a more formal level of representation. In this stage we are concerned with the organization of concepts into taxonomic hierarchies. We also aim to assert relationships and attributes at appropriate levels of abstraction within these hierarchies. The focus in terms of task modelling is to analyse tasks in terms of their task decomposition, knowledge requirements, information inputs and outputs and so forth. Typically, these modelling outcomes are an extension to earlier models detailed in task models and functional requirements specifications.

In the context of the current project we opted to use the modelling formalisms provided by the CommonKADS methodology (Schreiber et al, 2000). The CommonKADS methodology is generally consistent with other approaches, particularly frame-based approaches to knowledge infrastructure specification, and supports a rich variety of representational formats, e.g. UML notations.

The outcome of this stage in the knowledge engineering cycle is a set of knowledge models containing instances of the knowledge types described in Section 5.2. We adopt a standard delivery format for the communication and dissemination of these models to other project stakeholders.

5.3.3 Ontology Formalization

The final stage of the knowledge engineering cycle involves the formalization of Stage 2 knowledge structures. The formalization process entails the (partially) automatic translation of Stage 2 knowledge structures to a representational format that supports the requisite reasoning and decision-making capabilities of the prospective system within a semantic web environment. In this case we opted to use OWL, a W3C endorsed ontology language, as the representational basis for the knowledge infrastructure. OWL has a number of features that commend its adoption in the context of the current initiative, including:

- **RDF Compliance**: OWL is serialized in an RDF-compliant format and is therefore compatible with a range of semantic web technologies, including 3Store technology.
- **Tool Support**: A number of knowledge editing environments provide direct or indirect support for OWL, e.g. Protégé. In addition, some widely available APIs,

such as Jena, provide support for parsing, specifying and querying OWL repositories.

- **W3C Endorsement**: OWL is the ontology representation language endorsed by the W3C. This increases the likelihood that it will be adopted as an industry-standard format in the near future.
- **User Community**: OWL has gained widespread acceptance amongst the academic community and is currently being used in the context of a number of major research initiatives, e.g. MyGRID, AKT.
- **Semantic Expressivity**: OWL has a high level of semantic expressivity as a consequence of its description logic pedigree. Such expressivity lends support to certain forms of logico-deductive inference and subsumption reasoning.

The outcomes of this stage of the knowledge engineering cycle include an OWL ontology delivered either as an XML file or 3Store knowledge repository.

5.4 Domain Schemas

This section provides a provisional outline of the knowledge content for each of the domain schemas developed in the context of the current project[53].

5.4.1 Geography Ontology

This ontology deals with all the geographical aspects of the problem domain. It encompasses a wide variety of conceptualizations including terrain features, transport routes, rivers, shorelines, terrain elevation data, etc.

5.4.2 Transportation Ontology

This ontology covers all aspects of transportation. This may overlap, to some extent, with the geography ontology in the sense that transportation routes, e.g. airways and roads, may also be considered elements of the geographical (geo-spatial) domain.

5.4.3 Humanitarian Aid Ontology

This ontology covers information of relevance to humanitarian operations. It includes knowledge about humanitarian hazards (e.g. floods), humanitarian organizations, humanitarian aid programmes, humanitarian aid workers, and the types of resources that may be used for humanitarian relief operations.

5.4.4 Meteorology Ontology

This ontology deals with all aspects of the climate and weather. The meteorology ontology is important in enabling the prospective system to interpret and utilize information derived from local weather reports and forecasts as well as long term data about regional rainfall, snowfall, seasonal temperature, etc.

[53] Note that in this case we use the term domain schema and ontology synonymously. The idea is that the domain schema that is specified in the context of the Stage 2 Knowledge Models (see Section 5.3.2) will correspond to an OWL ontology at the Ontology Formalization stage (see Section 5.3.3).

5.4.5 Information Resources Ontology

This ontology details the information sources available to the prospective system. This includes the totality of information available from public domain databases, websites and web services, as well as briefings, emails and tactical datalink systems. It also includes a conception of the knowledge system itself which may serve as the source of internally-derived or inferred information (see Section 3.4)[54].

5.4.6 Geo-Political Ontology

This ontology details the conceptualizations used in the geo-political domain. This includes notions of countries, provinces, states, regions, settlements and the like. It is also subsumes ethnic and linguistic (perhaps also religious) groupings.

5.4.7 Military Ontology

This ontology includes all relevant conceptualizations in the military domain, including tactical operational areas and zones, military platforms, intelligence information, weapons, etc.

5.4.8 Datalink Ontology

This ontology details the information infrastructure of the tactical datalink systems used by the military to communicate information about the battlespace.

5.4.9 Equipment Ontology

This ontology details the various equipment items that may be used in the course of both military and humanitarian operations. It has substantial overlaps with the content of both humanitarian aid and military ontologies.

5.4.10 Knowledge System Ontology

This ontology details the problem-solving elements used by the knowledge system to fulfil its problem-solving objectives or to provide explanatory accounts of its own problem-solving and decision-making activities.

5.4.11 Agent Ontology

This ontology provides detailed characterizations of the various agents with which the system is required to inter-operate. The information captured by this ontology includes

[54] Two types of internal information are generated by the TDS: deliberative and reflective constructs. Reflective constructs are those constructs required to provide an explanatory account of reasoning activity or to justify decision-making outcomes to an external observer or inter-operating system. It includes a description of the conceptualizations used to represent the reasons for particular inferences in terms of knowledge-rich contingencies and information inputs. Such conceptualizations can be used to provide a (natural language) trace of problem–solving activity, which is useful for debugging, refinement and testing activities. Deliberative constructs include those constructs used by the prospective knowledge system to control or implement the requisite reasoning and decision-making competencies of the system. This type of construct includes standard problem-solving method constructs, such as classification, assessment and planning templates, as well as rules that express the knowledge-rich contingencies that inhere in the problem domain.

information about the operational role performed by the agent, the position of the agent in power and communication hierarchies, contact details associated with the agent, and information about the kinds of events the agent is subscribed to.

5.4.12 Communication Device Ontology

This ontology characterizes the various equipment items that are used to communicate or transfer information to inter-operating agents. The capabilities of a particular communication device are important in terms of limiting the kind of information that can be presented as well as the manner in which it should be presented (see Section 3.6) to end user agents (imagine the different constraints imposed by 19 inch colour monitor on the one hand and a palm-size PDA on the other).

5.5 Summary

This section has provided an overview of the strategy adopted for knowledge capture and representation in the context of the current initiative. It has also outlined the key areas of knowledge in which detailed knowledge modelling was undertaken and the kinds of problem-solving elements that will be used to realize the required reasoning and decision-making capabilities of the prospective system. The collection of knowledge models developed to describe the knowledge infrastructure of the target domain explicate in detail the kinds of concepts that can be instantiated by the system during the course of its reasoning activity. Such conceptualizations, in combination with rule contingencies, dictate the reasoning and inferential capability of the system and, ultimately, the kinds of decision outcomes it can yield.

6 System Evaluation

This section provides an overview of the evaluation metrics applied to the prospective system. Evaluation is important to ensure the system adequately addresses the original aims of the project, namely to improve situational awareness and operational effectiveness. A secondary goal is to ensure that the knowledge-related capabilities of the system are valid (e.g. in terms of making the correct decisions), reliable (capable of dealing with inconsistent or incomplete information, fault tolerant) and sufficiently generic (preservation of operational integrity across different situations and datasets). We can, and should, therefore, ask the following questions during the evaluation of the TDS:

- Is system performance preserved across different scenarios using different datasets? Are there specific aspects of the scenario that would suggest knowledge gaps or inadequacies in terms of system performance?
- How well do the decision outcomes made by the system match those expected by an end user, e.g. a domain expert?
- How well does the system cope with conflicting, incomplete or inconsistent data?
- To what extent does the system enhance situational awareness in human operatives? What impact does this have on operational performance?
- Does the system obey temporal constraints surrounding information provision, i.e. does it provide information in a timely manner that augments rather than hinders the problem solving process?

These questions serve as the basis for defining MOEs or MOPs that can be used to assess system performance.

With regard to the issue of validating the knowledge infrastructure (decision outcomes, fusion capabilities, etc.) it is common to develop simulation environments in which the desired knowledge capabilities can be tested with respect to specific scenarios. This simulation can be done in two ways:

1. **Paper-based simulation**: This method resembles a structured walk-through. The aim is to use the aforementioned scenario, and use the knowledge model to generate a paper trace of the scenario in terms of the status of knowledge model constructs and decision outcomes at each step of the scenario timeline. This can best be done in a table with three columns. The left column describes the steps in the scenario in terms of the terminology adopted for the domain. The middle column indicates how each step maps onto a knowledge-model element, e.g., an inference is executed with certain roles as input and output. The right column can be used for explanations and comments (see Figure 6-1).
2. **Simulation through a mock-up system**: This method relies on the development of a simulation environment for the system. Such an environment needs to have facilities for loading the knowledge-model specification plus a minimal set of implementation-specific pieces of code, such that the simulation can be done within a short time period (hours or days instead of weeks).

In the context of the current project we have opted for the latter of these two simulation options, i.e. the development of a mock-up system. The aim is to develop an early prototype of the TDS that can be used for both testing and presentational purposes.

Domain	Model	Explanation
The user says: "the car does not start".	DIAGNOSIS: **complaint**: engine-behavior.status = does-not-start	A complaint is received, for which a diagnostic task is started.
A possible cause is that the fuel tank is empty.	cover: hypothesis; fuel-tank.status = empty	One of the three possible causes is produced by this inference. The other two are "fuse blown" and "battery low".
In that case we would expect the gas indicator to be in the lowest regions.	predict expected-finding: gas-dial.value = zero	The expected finding provides us with a way of getting supporting evidence for this hypothesis.
System: "can you tell which value the gas dial indicates?". User: "It looks normal to me".	OBTAIN: actual-finding: gas-dial.value = normal	This is not what we expected, so we can rule out this possible fault.
The values differ, so it cannot be an empty fuel tank..	match: **result** = not-equal	The test to find supporting evidence fails.
We go and look for another possibility.	cover: **hypothesis** = battery.status == low	We repeat the process with a second possible solution.
And so on	.....	The task-control loop continues.

Figure 6-1:Paper Simulation of Knowledge System Activity

7 Collaborative Opportunities

This section provides an overview of potential opportunities for collaboration and technology transfer within the DIF DTC. At a recent DIF DTC Theme Meeting[55], held on the 28th October 2004, Martin Ferry, head of the DIF DTC Thematic area entitled Situational Awareness and Human Factors, emphasized the need for collaboration between DTC project members, especially with respect to the forthcoming round of DTC submissions for Phase II funding. Preference will be given to Phase II proposals that emphasize collaboration among the DTC initiatives, particularly with respect to mutual exploitation of technological and intellectual artefacts. In order to assess the opportunities for collaboration within the DIF DTC it is important to understand the kind of initiatives currently being undertaken. This section aims to detail each of the thematic areas (see Section 7.3) and projects (see Section 7.4) within the DIF DTC as the basis for subsequent technology exploitation and assessments of collaborative potential[56]. Section 7.1 provides a general introduction to the DTC, while Section 7.2 presents a schema for knowledge capture in this area.

7.1 DTC Overview

The DTCs[57] were established by the MOD to undertake leading edge research within a number of problem areas in order to yield novel technologies with defence- and government-related applications. In the first Phase (Phase I) three DTCs were established:

- **Data and Information Fusion (DIF)**: the aim of this consortium is to explore ways in which data and information from a variety of different sources can be fused to improve situational awareness and enhance operational effectiveness throughout the chain of command. The consortium is led by General Dynamics UK Ltd. Other members of the consortium are BT Exact, QinetiQ, Imperial College and the Universities of Bristol, Cardiff, Cambridge, Southampton, de Montford, Surrey and Cranfield.
- **Human Factors Integration (HFI)**: this consortium investigates the relationship between human factors and a range of defence capabilities. The consortium is led by Aerosystems International. Other members are Birmingham University, Brunel University, Cranfield University, Lockheed Martin UK Integrated Systems Ltd, MDBA Ltd, Systems Engineering and Assessment Ltd and VP Defence Ltd.
- **Electromagnetic Remote Sensing (ERS)**: this consortium aims to undertake a broad programme of research aimed at enhancing the performance of sensing equipment in a cost effective manner. This consortium is led by BAES (Edinburgh). Other members include Thales Defence UK Ltd, Roke Manor

[55] see DTC/Notes-28-10-2004#1

[56] Most of this information was gathered during the DIF DTC conference held on 17th May 2004 at the JSCSC, Shrivenham.

[57] http://www.mod.uk/dtc/index.html

Research Ltd and Filtronic Plc. QinetiQ, together with a number of universities and small and medium enterprises are sub-contractors in the consortium.

Projects in each of the DTCs are subdivided into a number of thematic areas based on the research focus of the projects. These themes are detailed in Section 7.3 (for the DIF DTC). The DTCs represent a consortia of major players from industry and academia, with projects partly funded by the MOD and industrial collaborators. General Dynamics UK Ltd. serves as the principal industrial contractor and funding authority for the DIF DTC.

7.2 Information Schema

This section presents a schematic overview of the DTCs and constituent projects as the basis for assessments of each project in terms of its potential value and relevance to the current DTC initiative. The schema details the relation between DTCs, thematic areas, projects and project participants and could be used as the basis for database development initiatives geared to maintaining a repository of information about other DTC projects.

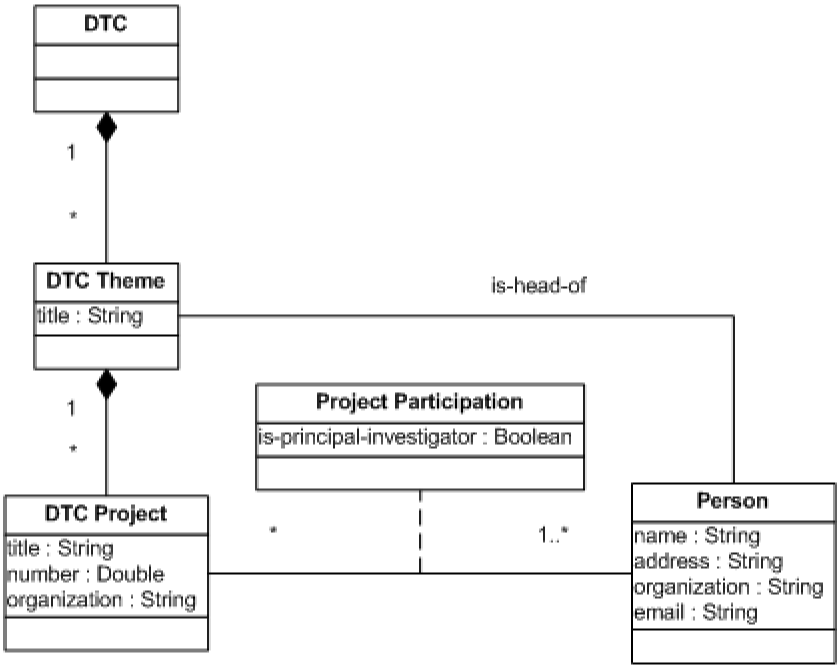

Figure 7-1:DTC Schema

Figure 7-1 illustrates a partial schema for DTC project information. Each DTC is composed of a number of thematic areas, each of which is itself composed of projects. Each project is associated with a number of project participants who may serve as principal investigator for the project. As was mentioned above, this information is easily transformed into a database specification for monitoring project progress and storing contact information about project participants.

7.3 DTC Themes

Each DTC is divided into a number of thematic areas that group related research initiatives. This section details the thematic areas for the DIF DTC. The descriptions provided for each of the following thematic areas are adapted from the DIF DTC website[58]. The name in brackets after the section header indicates the theme leader for the thematic area.

7.3.1 Agents and DIF Architectures (Dr Robert Ghanea-Hercock)

This theme considers the trade off between application dependent or bespoke DIF architectures and standards, as traditionally used in military applications, compared to fixed architectures including COTS, and/or self organising architectures and standards for data fusion allied with appropriate performance metrics that are platform independent. This theme considers the affects of C4ISTAR DIF systems via a unified modelling language. Adaptive Agent techniques will enable distributed, diverse, uncertain and conflicting data to be managed and visualised. Other approaches including Neurofuzzy, Bayesian, and Case Based Reasoning etc. will be benchmarked with agent methods to evaluate the most parsimonious DIF architectures for ISTAR sensors in decision support.

In Phase II of the DIF DTC this theme will aim to address the following basic research areas of knowledge representation, systems integration, architectures, complexity, knowledge discovery and high level abstraction.

7.3.2 Multi-Dimensional Fusion (Professor Dave Bull)

Conventionally data fusion has focussed on 2D sensory sources (e.g. range or range rate and time) however many data sources (video, imagery etc.) have higher dimensionality. This theme will endeavour to investigate data fusion from multi-dimensional data sources. The theme also aims to fully exploit ISTAR sensor potential in terms of 2D-4D integration (conventional spatial coordinates/variables and time) as well as multi-dimensional fusion via multi-and-hyperspectral image techniques.

7.3.3 Multi-Sensor Management (Dr Maria Petrou)

This theme focuses on the issue of optimal sensor coverage resource management and control via DIF feedback, as well as sensor conditioning and intelligent sensor management, active data gathering and data mining, distributed sensing and communications/bandwidth management.

7.3.4 Situation Awareness and Human Factors (Martin Ferry)

This thematic area, of which the current initiative is a part, aims to investigate the impact of data fusion on situation assessment and situational awareness. The theme addresses issues such as the uncertainty associated with data quality and fusion processes. The theme also incorporates human factors research. The human component within a DIF system is a critical component, not only as a data integrator but also as a decision maker. This theme considers human performance as part of high-level sensor fusion (including improved interaction), and human orientated optimised picture compilation from multiple data sources and optimised information displays.

[58] http://www.difdtc.com

7.3.5 Tracking and Target Classification (Dr Simon Godsill)

This theme extends the conventional DIF approach of Kalman filtering based on linear process models in clutter free domains to data based, non-linear stochastic multiple targets in clutter with high precision in both tracking and target identification.

7.4 DTC Projects

This section describes the projects undertaken in the context of the DIF DTC within the various thematic areas outlined in Section 7.3. Each project is described in terms of:

- **Title**: the title of the project
- **Description**: a description of the research agenda of the project
- **Project Number**: the number assigned to the project within the DTC
- **Participants**: the names of individuals who participate in the project
- **Organization**: the organization responsible for undertaking the research

7.4.1 Multi-Sensor Management Projects

These projects form part of the Multi-Sensor Management theme (see Section 7.3.3)

Title:	Communication Optimisation for Distributed Sensor Systems		
Number:	2.2	**Organization:**	University of Bristol
Description:	This project aims to optimize the communication profiles of distributed sensor systems so as to improve information fusion and situational awareness. The emphasis here is on how information should be propagated between a collection of distributed sensor devices in order to maximize information throughput to executive agencies. Key issues to address in this respect include adaptive network topologies and self-configuration of sensor devices to account for environmental conditions, particularly those that impact on RF propagation. The project is additionally investigating the issue of image fusion over sensor networks (e.g. images from visible light and infra-red cameras).		
Participants:	Prof David Bull, Prof Andrew Nix		

Title:	Characterisation, Modelling and Mitigation of Impairments produced by Radio Channels, Sensors & Communication Equipment for Data Fusion		
Number:	4.7	**Organization:**	University of Cardiff
Description:	This project was not discussed at the DTC conference talk given by Dr Maria Petrou, but it seems to involve the use of empirical methods to build a library of sensor fault signatures that can be used to improve target tracking in cluttered and hostile environments. The projects appears to involve the modelling of communication channels and faults.		
Participants:			

Title:	Condition Monitoring and Fault Detection		
Number:	6.3	**Organization:**	Imperial College, London
Description:	This project was not discussed at the DTC conference talk given by Dr Maria Petrou, but it seems to involve the use of analytical methods to build a library of sensor fault signatures that can be used to improve target tracking in cluttered and hostile environments. The emphasis seems to be the analysis of fusion error scenarios to develop fault detection filters based on a library of fault signatures.		
Participants:			

Title:	Distributed Data Sensors with Low Communications Bandwidth		
Number:	6.5	**Organization:**	Imperial College, London
Description:	This project aims to investigate the use of large numbers of small, self-contained sensor packages (nodes) to gather information within an area of interest. The nodes are randomly distributed and automatically form wireless arrays, which improves fault tolerance and network adaptability.		
Participants:	Dr A Manikas, Miss Cara Beck, George Elissaios (PhD Student)		

Title:	Modelling the RF Environment for Single and Multi-Sensor Data Fusion Applications		
Number:	7.1	**Organization:**	QinetiQ
Description:	This project aims to investigate the RF propagation environment for optimizing information transmission across RF communication channels. Local environmental conditions and terrain features can profoundly affect the propagation of RF radiation. The intelligent incorporation of these factors into the fusion process may serve to minimise the number of fusion errors.		
Participants:	Dr Anil Shukla		

Title:	Latent Fault Detection in Large-scale Sensor Networks		
Number:	7.3	**Organization:**	QinetiQ
Description:	The aim of this project is to develop procedures to detect and classify sensor faults in large distributed sensor arrays. The hope is that health monitoring of sensor networks and their components will yield improvements in the cost-effective use of sensor equipment and battlefield surveillance.		
Participants:	Dr Richard Glendinning		

Title:	Multi-Sensor Active Management		
Number:	8.1	**Organization:**	University of Southampton
Description:	The aim of this project is to manage sensor devices so as to optimize information gathering activities and improve situational awareness. The key goal is to select and orient the right sensors to the right object (target) at the right time.		
Participants:	Prof Neil White, Prof Chris Harris, Dr Alexander Dolia (PostDoc), Scott Page (PhD student)		

Title:	Optimal Signal Extraction and Sensor Modelling Algorithms		
Number:	8.5	**Organization:**	University of Southampton
Description:	The aim of this project is to design intelligent sensors that perform local fault diagnosis for the purposes of uncertainty measurement. Intelligent sensor devices will improve uncertainty estimates about the information gathered and thereby improve information fusion based on those measures.		
Participants:	Prof Neil White, Prof Chris Harris, Dr Tien Tran		

Title:	Sensor Clustering for Robust Control Systems		
Number:	8.9	**Organization:**	University of Southampton
Description:	The aim of this project is to extend the capabilities of extant active noise cancellation systems used by some civilian aircraft by investigating the role of distributed arrays of structural actuators and sensors as opposed to fully centralized control systems.		
Participants:	Prof Steve Elliot		

7.4.2 Tracking and Target Classification Projects

These projects form part of the Tracking and Target Classification theme (see Section 7.3.5)

Title:	High Precision Bearings Only Tracking for Manoeuvring Targets		
Number:	6.1	**Organization:**	Imperial College, London
Description:	This investigates difficult tracking problems where generic tracking methods, such as Kalman filtering, multiple linear model techniques, particle filtering, etc) break down or give poor results. The aims are to reveal the pathologies inherent in these problems, provide bounds on performance of ideal trackers and to develop new tracking algorithms that exploit most efficiently such information about target motion as is available.		
Participants:	J.M.C Clark, R.B. Vinter, M.M Yaqoob		

Title:	Adaptive Markov Models for Signature Extraction from Multiple Sensors		
Number:	6.6	**Organization:**	Imperial College, London
Description:	This project aims to investigate the use of Hidden Markov Models (HMMs) for improving the data-driven modelling of time-varying target signatures. The key objective seems to be use of Markov modelling techniques to describe the dynamic behaviour of target objects. The hope is that such information will yield improvements in target identification and classification.		
Participants:	Mike Brookes		

Title:	Future Data Fusion Systems		
Number:	6.8	**Organization:**	Imperial College, London
Description:	This project has developed a basic mathematical framework, based on Multi-dimensional Infinite State Space Continuous Time Markov Chains (ISSCTMC) and Markov Regenerative Processes (MRP), to improve target identification. The model assumes that target objects have intentions and relationships to other objects, which will govern their behavioural dynamics interaction.		
Participants:	Erol Gelenbe, Varol Kaptan, Yu Wang		

Title:	Model Based Tracking		
Number:	7.2	**Organization:**	QinetiQ
Description:	This project aims to build a comprehensive repository of tracking algorithms as the basis for algorithm assessment.		
Participants:	S Maskell		

Title:	Target Classifier System		
Number:	7.7	**Organization:**	QinetiQ
Description:	This project focuses on the development of a hybrid classification system for airborne targets that exploits motion information and information derived from multiple identity/attribute sensors.		
Participants:	Keith Copsey		

Title:	Biometric Fusion Methods for Human Recognition in Secure Environments		
Number:	8.10	**Organization:**	University of Southampton
Description:	The aim of this project is to integrate morphometric facial feature (earlobe) information derived with gait structure information as the basis for biometric identification.		
Participants:	Mark Nixon, John Carter		

Title:	Multi-sensor Tracking and Classification		
Number:	10.2	**Organization:**	University of Cambridge
Description:	This projects aims to improve multiple target tracking in high clutter environments using a number of statistical and machine learning algorithms.		
Participants:	Simon Godsill, Arnaud Doucet		

7.4.3 Multi-Dimensional Fusion Projects

These projects form part of the Multi-Dimensional Fusion theme (see Section 7.3.2)

Title:	Image and Video Sensor Fusion		
Number:	2.1	**Organization:**	University of Bristol
Description:	The aim of this project is to develop a framework for image and video fusion that is robust in the face of fusion challenges, including low light levels, noise, distributed sensor arrays, multiple resolutions and multiple spectra (e.g. visible, IR, etc)		
Participants:	David Bull, Guy Nason, Nishan Canagarajah, Stavri Nikolov, Artur Loza, Alessandro Cardinali, Tim Dixon, John Lewis		

Title:	Statistical Data Fusion of Battlefield Data		
Number:	4.10	**Organization:**	University of Cardiff
Description:	The aim of this project is to develop robust fusion algorithms for digitized battlespace information, e.g. information transmitted over tactical datalinks.		
Participants:	David Marshall, Gavin Powell		

Title:	Object-Based Hyper-spectral Image Fusion		
Number:	6.4	**Organization:**	Imperial College, London
Description:	The aim of this project is to develop algorithms for multi-spectral images and using image enhancement and boundary-based segmentation techniques.		
Participants:	Tania Stathaki, Qi Li, Nikolaos Mitianoudis		

Title:	Hyper-spectral Data Fusion with Broadband FLIR for Air-Ground Target Tracking		
Number:	7.11	**Organization:**	QinetiQ
Description:	This project aims to investigate the fusion of information from broadband FLIR/IRST and hyperspectral sensors. The hope is that by co-opting the relative capabilities of these sensors with respect to a number of performance parameters (e.g. spatial resolution, spectral resolution, signal-to-noise-ratio, etc.) the detection and identification of long range targets can be improved.		
Participants:	Mike Nicholas, Joanne Nothard		

Title:	Adaptive Hyper and Multi-Spectral Data Fusion for Target Detection and Tracking, Biometric Identification		
Number:	8.2	**Organization:**	University of Southampton
Description:	The aim of this project is to extend existing data fusion capabilities in the areas of target tracking and speaker recognition.		
Participants:	Prof Bob Damper, Dr Steve Gunn, Prof Chris Harris, Dr Baofeng Gao		

Title:	Sub-Pixel Target Detection		
Number:	9.3	**Organization:**	University of Surrey
Description:	The aim of this project is to combine (fuse) a number of low resolution images, with sub-pixel registration alignment errors, to produce a single high resolution image. A second objective of this project is to multispectral images to detect the presence of certain sub-pixel targets in an image.		
Participants:	Maria Petrou, Christos Papathanassiou		

Title:	Point and Image Data Fusion for Target Signature Detection in Clutter		
Number:	9.7	**Organization:**	University of Surrey
Description:	The aim of this project is to demonstrate improved decision-making based on fused information from satellite sources. The focus areas for improved decision-making include the passive detection and tracking of ships and the classification of ship type from wake information.		
Participants:	Phil Palmer, Stephen Mackin, Yiping Sun		

Title:	Multi-resolution Image Retrieval		
Number:	10.1	**Organization:**	University of Cambridge
Description:	This project aims to develop tools to automate the analysis and classification of battlefield images. A key technical thrust for this project is the development of probability models to identify if events have occurred in a scene based on large sets of images captured from different sensors at different times.		
Participants:	Nick Kingsbury, Julien Fauqueur, Ryan Anderson		

7.4.4 Agents & DIF Architectures Projects

These projects form part of the Agents & DIF Architectures theme (see Section 7.3.1)

Title:	Agent Managed Sensor Networks		
Number:	1.1	**Organization:**	British Telecom
Description:	The aim of this project is to facilitate NEC via autonomous ICT systems.		
Participants:	Robert Ghanea-Hercock, Nima Kaveh		

Title:	Agents Supporting a Decision Desktop		
Number:	7.6	**Organization:**	QinetiQ
Description:	This project aims to develop a software tool for improved information fusion and decision support. The application features highly customizable multiple visualizations of information collected from a variety of sources (see Section 4.8.4)		
Participants:	David Allsop		

Title:	Design and Implementation of an Agent-based Control Systems for Distributed Data Fusion		
Number:	8.6	**Organization:**	University of Southampton
Description:	This project aims to develop flexible and robust methods for managed decentralized processes in a logistics supply chain management scenario.		
Participants:	Professor Nick Jennings, Dr Xudong Luo, Dr Esther David, Perukrishnen Vytelingum, Talal Rahwan		

Title:	Decision Support System Using Dynamic Parsimonious Information Fusion Architectures		
Number:	9.1	**Organization:**	University of Surrey
Description:	The aim of this project is to develop a decision support system capable of intelligent information fusion and presentation in order to provide an improved C4I/ISTAR capability. The project is exploiting Bayesian belief networks, case-based reasoning and CFS technology to provide an intelligent basis for improved situational awareness via information fusion.		
Participants:	Panos Louvieris, Bob O'Keefe, Jan Powell-Perry, Gareth White Nataša Mašanović, Chun-Quan Tang, Ying Zhang		

Title:	Self-Organising Network of Networks		
Number:	11.1	**Organization:**	General Dynamics UK Ltd.
Description:	This project aims to exploit artificial intelligence to create self-sustaining military network systems given a set of rules and guidelines. To key focus here is to develop software agents that are able to engage in a number of self-directed activities, including configuration, optimization, protection and repair. Analogies are made with real-world biological components that participate in these activities to preserve and maintain the functional integrity of some larger system.		
Participants:	Graham Atkins, John Salt, Simon Robinson, James Wise, James Spillings, Greg Phillips, Peter Burke, Gareth Smith		

7.4.5 Situation Awareness and Human Factors Projects

These projects form part of the Situation Awareness and Human Factors theme (see Section 7.3.4)

Title:	Human Computer Interaction for DIF		
Number:	1.2	**Organization:**	British Telecom
Description:	This project focuses on the use of haptic/tactile displays to reduce cognitive workload and efficiently distribute the allocation of sensory and cognitive resources with respect to militarily-relevant information processing tasks.		
Participants:			

Title:	Mixed Reality System for Urban Environments		
Number:	1.15	**Organization:**	British Telecom
Description:	This project aims to develop mixed-reality systems for urban combat environments as a means to improve situational awareness and operational effectiveness.		
Participants:	Andrew Gower, Barry Crabtree, James Bulman, Alex Loffler, Matthew Polaine, Jon Sutton, Dale Robertson, Matthew Iles, Martin Trimby		

Title:	Designing Integrated Displays to Support Team Situation Awareness		
Number:	4.9	**Organization:**	University of Cardiff
Description:	This project aims to experimentally evaluate the putative benefits of information fusion with respect to situational awareness using a variety of cognitive performance measures.		
Participants:			

Title:	Integrated Text Analysis		
Number:	7.5	**Organization:**	QinetiQ
Description:	The aim of this project is automatically process the contents of textual sources as a means of increasing situational awareness. The group has evaluated a number of technologies, such as inductive logic programming (ILP), machine learning (ML) and natural language processing (NLP), for classifying the text contents of documents with respect to semantically relevant categorizations, e.g. recognising that a text resource is describing a type of aircraft.		
Participants:	Dr Claire Thie		

Title:	Novel HMI Concepts		
Number:	7.8	**Organization:**	QinetiQ
Description:	This project focuses on the effort to develop systems for automatic gesture recognition. The current focus of the project is on hand gestures.		
Participants:			

Title:	Data Fusion Sensor Networks for Well-being Monitoring Applications		
Number:	1.7	**Organization:**	British Telecom
Description:	The aim of this project is to monitor aspects of task performance or cognition and undertake some action if desired. An example could be to raise an alert if the number of typing errors on a data entry task falls below a threshold level. A key feature of this project is the use of knowledge to infer when some form of intervention (and what kind of intervention) is required given some significant change in cognition or task performance.		
Participants:			

7.5 Summary

This section has also provided an overview of the potential basis for collaborative activities within the DIF DTC. A number of research areas have been identified that could serve as the basis for future collaborative efforts or joint proposals. These include:

- **HMI Technology**: the development of interfaces that support the 'effective' visualization and processing of battlespace information so as to improve the problem-solving competency and operational effectiveness of military personnel.
- **Augmented Cognition**: the exploitation of key technologies that aim to extend an agent's cognitive profile via computational technologies that are explicitly designed to address bottlenecks, limitations, and biases in cognition.
- **Decision Support**: the use of knowledge technologies to assist problem-solving agents and knowledge workers with strategic decision-making activities.
- **Situational Awareness**: improved situational awareness to establish a strategic advantage with respect to operational objectives.

Within each of these research areas we have attempted to identify specific topics for further research that relate to the research agendas of extant projects within the DIF DTC. They include:

- **HMI Technology**:
  - semantically-enriched characterizations of display device capabilities
  - optimal display of information vis-à-vis display device characteristics
  - optimal display of information vis-à-vis operator role and problem-solving profile
- **Augmented Cognition**:
  - monitoring of operator capabilities and cognitive state
  - knowledge-based selection of presentation modality based on operator state and workload
- **Decision Support**:
  - adaptive sensor configuration for optimum signal detection and information transmission
  - knowledge-based configuration of sensor network topologies

  - exploitation of semantically-enriched information for decision support
  - meta-level semantic characterizations of intelligent sensors for information fusion
- **Situational Awareness**:
  - knowledge-based filtering and selective attention to contextually-relevant information from other sensors and military platforms
  - semiometric information harvesting and selective attention to mission critical information

In all likelihood this provisional list of topic areas will be refined and extended following further contact with potential DIF DTC collaborators and a review of the MODs current and prospective capability requirements.

8 Exploitation & Dissemination

The DIF DTC does not exist solely to undertake pure research; it is also described as a 'Centre of Excellence'[59] whose aims include the exploitation of research results in both civil and defence contexts. To this end we have identified a number of exploitation routes for the intellectual and technological deliverables of the current project.

8.1 Conferences

An abstract has been submitted for the 8th International Conference on Information Fusion[60] to be held in Philadelphia, June 2005. Other candidate conferences include:

- 4th International Semantic Web Conference (ISWC)[61]
- 19th International Joint Conference on Artificial Intelligence (IJCAI) [62]
- 14th International World Wide Web Conference (WWW2005)[63]
- 15th International Conference on Knowledge Engineering and Knowledge Management (EKAW)
- 2nd European Semantic Web Conference (ESWS)[64]
- International Conference on Web Services (ICWS)[65]
- 11th International Conference on Human-Computer Interaction[66]
- 20th National Conference on Artificial Intelligence (AAAI)[67]
- 3rd IEEE International Conference on Intelligent Systems[68]
- European Conference on Artificial Intelligence (ECAI)[69]

In addition to conferences we also aim to present the results of our research at invited talks and presentations that highlight the technical expertise of the IAM group within the University of Southampton.

---

[59] Dr Andy TilBrook – DIF DTC Conference, JSCSC Shrivenham

[60] http://www.fusion2005.org/

[61] http://iswc2005.semanticweb.org/

[62] http://www.ijcai-05.org/

[63] http://www2005.org/

[64] http://www.eswc2005.org/

[65] http://conferences.computer.org/icws/2005/

[66] http://www.hci-international.org/

[67] http://www.aaai.org/Conferences/National/2005/aaai05.html

[68] http://www.ieee.org.uk/docs/is2006.pdf

[69] http://www.dsic.upv.es/ecai2004/

8.2 Publications

It is important that the research outcomes of the current initiative are disseminated to a wide academic audience. For this reason we aim to publish the results of our research and development efforts in high quality journals such as IEEE Intelligent Systems and other journals within the AI and Semantic Web community.

8.3 Website

The web can serve as a valuable means of advertising the research aims and objectives of the current initiative to a wide range of interested parties. As such we aim to produce a website to promote the current project, either independently or as part of the AKT initiative. Ideally, the website should enable interested users and organizations to register their interest so that they can be informed of new developments and research outcomes via regular newsletters. Such information will also provide us with a list of contacts for further exploitation and dissemination activities.

8.4 AKT Integration

The AKT initiative within the University of Southampton has well established exploitation and dissemination routes that can be easily exploited by the current initiative. Exploitation routes include AKT town meetings, the AKT website and contact with AKT partners/stakeholders.

8.5 Civilian Applications

Our work in the context of the current project is of relevance to a number of applications in the civilian sector. The focus on humanitarian operations suggests that any number of humanitarian aid agencies may be interested in the outcomes of the current project, particularly those aid agencies that are currently involved in, or that are likely to be involved in, relief efforts undertaken against a backdrop of civil unrest and military activity. More generally, the results of our work in this area will be of interest to any organization in which successful knowledge actions (e.g. decisions about the disposition of organizational assets) depend on improved situational awareness. For example, we suggest that the technological deliverables of the project could be easily adapted to meet the requirements of agencies within the emergency services, e.g. fire, police, ambulance, etc.

9 Deliverables Roadmap

This section describes the deliverables to be delivered in the context of the current project[70]. Each deliverable represents the intermediate results or final outcome of research and development activities undertaken with respect to the project aims and objectives described in Section 1.1. Section 9.1 outlines the key areas of research and development in the form of a series of work packages. Section 9.2 describes the deliverable items to be delivered in each of these work packages.

9.1 Work Packages

The programme of work to be undertaken in respect of the current project can be divided into a number of parallel activity areas, e.g. knowledge modelling, project administration, TDS development, etc. In this section we outline these activity areas in the form of a number of work packages. The deliverables to be produced in each of these work packages is detailed in Section 9.2.

9.1.1 WP100 Scenario Specification

This work package focuses on the specification of a humanitarian aid scenario. The scenario provides the basis for testing and validation activities (see Section 2.1) that will facilitate the knowledge engineering initiatives undertaken in the context of WP200. In addition, the scenario provides the basis for most of the experimental evaluation activities described in Section 6 and also serves as a presentational device, which can be used to demonstrate the capabilities of the prospective TDS (to be delivered in WP300) to a variety of project stakeholders and external agencies.

9.1.2 WP200 Knowledge Engineering

This work package subsumes all the knowledge engineering activities undertaken in the context of the current project. It focuses on the specification of the knowledge infrastructure (see Section 5) of the prospective system. Key activities undertaken in the context of this work package include knowledge acquisition and modelling, ontology formalization and the production of an instantiated knowledge repository, i.e. 3Store.

9.1.3 WP300 System Implementation

The work undertaken in the context of this work package includes all the system development activities, i.e. implementation of the prospective TDS and its prototypes. The implementation activity can itself be sub-divided into a number of distinct initiatives, each of which focuses on a particular aspect of the TDS. These include implementation of the knowledge system components, system architecture specification and interface design.

[70] The description of project deliverables in this section overrides an earlier description provided in the Scenario Specification report (ref: DTC/WP100/Scenario). This earlier description should now be ignored in favour of the current description.

9.1.4 WP400 Exploitation & Dissemination

This work package focuses on the establishment of exploitation and dissemination routes as discussed in Section 8. It includes efforts aimed at establishing contacts with other DIF DTC groups, website implementation, conference attendance and paper publications.

9.1.5 WP500 Project Management & Administration

This work package covers all activities related to the management and administration of the project. It includes the design of review materials, authoring of technical reports, attendance at DTC review meetings and conferences, production of project management materials and general project administration.

9.2 Deliverables

This section describes the deliverables to be produced in each of the work packages described in Section 9.1. Each deliverable is described in terms of the following:

- **Title**: the title of the deliverable item
- **Description**: a brief description of the deliverable
- **WP ID**: the unique work package identifier assigned to the deliverable
- **Reference**: the reference number assigned to the deliverable, if any
- **Type**: the type of the deliverable, e.g. software, documentation, etc.
- **Due Date**: the expected completion date for the deliverable
- **Location**: the URI of the deliverable in relation to the project's root directory

9.2.1 WP100 Deliverables

Title:	Scenario Specification Document		
WP ID:	WP100	**Reference:**	DTC/WP100/Scenario
Description:	Provides a provisional characterization of the problem domain and scenario. The scenario is presented in a narrative format accompanied by a series of scenario timelines.		
Due Date:	31/03/2004	**Type:**	Document
Location:	DTC/WP100/Scenario Specification		

Title:	Scenario Instantiation Toolkit		
WP ID:	WP110	**Reference:**	NA
Description:	A software application geared towards the detailed instantiation of scenario events and information as described in the earlier Scenario Specification Document. The tool should provide facilities for editing and visualizing scenarios using a graphical interface.		
Due Date:	31/03/2005	**Type:**	Software
Location:	DTC/WP100/Scenario Instantiation Toolkit		

Title:	Instantiated Scenario Suite		
WP ID:	WP120	**Reference:**	NA
Description:	Provides a detailed specification of scenarios developed with the aforementioned Scenario Instantiation Toolkit. The scenarios should be delivered as a series of XML files.		
Due Date:	31/05/2005	**Type:**	XML File
Location:	DTC/WP100/Scenarios/		

9.2.2 WP200 Deliverables

Title:	Knowledge Model Document		
WP ID:	WP200	**Reference:**	DTC/WP200/Knowledge
Description:	Details the results of knowledge modelling in the target domain. The Knowledge Model Document details the problem-solving types and structures that are part of the knowledge infrastructure of the proposed TDS. The structure of the knowledge document follows the recommended structure for communication of the CommonKADS Knowledge Model (see Schreiber et al, 2000).		
Due Date:	31/12/2004	**Type:**	Document
Location:	DTC/WP200/Knowledge Model/		

Title:	Knowledge Web		
WP ID:	WP210	**Reference:**	NA
Description:	Details the results of knowledge modelling in the target domain. In this case the knowledge model contents are presented as a knowledge web to facilitate knowledge infrastructure navigation and information retrieval.		
Due Date:	31/12/2004	**Type:**	Website
Location:	DTC/WP200/Knowledge Web/		

Title:	Task Model Document		
WP ID:	WP220	**Reference:**	DTC/WP200/Task
Description:	Details the task model developed for the current initiative. The Task Model analyses the tasks to be implemented by the prospective system. The structure of the task model document follows the recommended structure for communication of the CommonKADS Task Model (see Schreiber et al, 2000).		
Due Date:	28/02/2005	**Type:**	Document
Location:	DTC/WP200/Task Model/		

Title:	Knowledge Repository		
WP ID:	WP230	**Reference:**	NA
Description:	Specifies the knowledge infrastructure of the prospective system in the form of a populated 3Store. The 3Store provides a mechanism to store domain ontologies described using OWL. It also supports interfaces to query and retrieve information from the repository.		
Due Date:	31/12/2004	**Type:**	Database
Location:	DTC/WP200/3Store/		

Title:	OWL Ontology		
WP ID:	WP240	**Reference:**	NA
Description:	Specifies the OWL ontology for the current problem domain or domain of discourse. The OWL ontology corresponds to a logically formalized representation of the knowledge structures described in the CommonKADS knowledge models.		
Due Date:	31/12/2004	**Type:**	Document
Location:	DTC/WP200/Ontology/		

9.2.3 WP300 Deliverables

Title:	TDS Prototype		
WP ID:	WP300	**Reference:**	NA
Description:	Represents a prototype of the prospective system used to demonstrate system capabilities and test the integrity of knowledge processes.		
Due Date:	28/02/2005	**Type:**	Software
Location:	DTC/WP300/Prototype/		

Title:	TDS		
WP ID:	WP310	**Reference:**	NA
Description:	Represents the final software artefact to be delivered by the project, i.e. a technical demonstrator system to showcase the role played by knowledge and semantic web technologies in improving situational awareness.		
Due Date:	31/10/2006	**Type:**	Software
Location:	DTC/WP300/TDS/		

9.2.4 WP400 Deliverables

Title:	Website		
WP ID:	WP400	**Reference:**	NA
Description:	Represents the website used to promote the activities of the research group, register the interest of interested parties and disseminate information via newsletters.		
Due Date:	30/04/2005	**Type:**	Website
Location:	DTC/WP400/Website/		

Title:	Exploitation and Dissemination Review Document		
WP ID:	WP410	**Reference:**	DTC/WP410/Review
Description:	Provides a review of exploitation and dissemination activities undertaken up to 30th June 2005.		
Due Date:	30/06/2005	**Type:**	Document
Location:	DTC/WP400/Review/		

Title:	Exploitation and Dissemination Review Document		
WP ID:	WP420	**Reference:**	DTC/WP420/Review
Description:	Provides a review of exploitation and dissemination activities undertaken up to 30th June 2006.		
Due Date:	30/06/2006	**Type:**	Document
Location:	DTC/WP400/Review/		

9.2.5 WP500 Deliverables

Title:	Technical Progress Report		
WP ID:	WP500	**Reference:**	DTC/WP500/Progress
Description:	Provides a technical progress in the project up to 30th November 2004.		
Due Date:	30/11/2004	**Type:**	Document
Location:	DTC/WP500/Technical Progress/		

Title:	Technical Progress Report		
WP ID:	WP510	**Reference:**	DTC/WP510/Progress
Description:	Provides a technical progress in the project up to 30th November 2005.		
Due Date:	30/11/2005	**Type:**	Document
Location:	DTC/WP500/Technical Progress/		

Title:	Quarterly Progress Review Report		
WP ID:	WP520	**Reference:**	NA
Description:	A series of reports are to be submitted at quarterly intervals throughout the project lifecycle. Reports should be submitted to the relevant funding authority at the end of each of the following months: • February • May • August • November		
Due Date:	NA	**Type:**	Document
Location:	DTC/WP500/QPR/		

10 Feasibility Issues

This section presents the results of a feasibility analysis based on the current project status. The feasibility analysis is important for the identification of risks and uncertainties that may jeopardize future project progress. The early identification and resolution of these concerns is a key aspect in ensuring a successful outcome to the project.

We have undertaken a comprehensive feasibility analysis based on the following feasibility areas:

- **Business Feasibility**: addresses the benefits versus the costs of the proposed solution. Business feasibility assessments focus on the identification of the costs and benefits associated with introducing a knowledge system solution.
- **Project Feasibility**: addresses the further project actions that need to be undertaken to ensure the adequate realization of project goals and objectives.
- **Technical Feasibility**: addresses concerns about whether the needed technologies for the solution are available and within reach.

Subsequent sections detail the specific issues and concerns raised with respect to these areas by responding to a number of feasibility criteria.

10.1 Business Feasibility

What are the expected benefits for the organizational stakeholders involved in the project? (BF1)

We believe the potential benefits to be delivered by this project are considerable. The project aims to build on current efforts to improve situational awareness with respect to humanitarian operations and the activities of military personnel. In addition, the project aims to avoid a situation of information overload by using semantic annotations to focus attention on information of direct relevance for current operational objectives and epistemic concerns. The use of semantically-enriched information to implement a selective attentional filter in this context is relatively new and we believe its full potential is yet to be realized.

Direct customer benefits aside, the current project aims to further extend the exploitation of semantic web technology and thus serves as a useful building block in the evolution of intelligent systems, the semantic web and semantic web services.

How relevant is the prospective system to the capability requirements of stakeholder groups? (BF2)

The current work is highly relevant to the capability requirements of the UK defence agencies. Recent, current and, in all likelihood, future conflicts, requiring the participation of UK defence forces, have been conducted either for humanitarian purposes or have directly occurred against a backdrop of humanitarian intervention. Cooperation with humanitarian agencies and improved awareness of the temporal unfolding of humanitarian events is therefore critical to the success of military forces in conflict situations. 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.

10.2 Project Feasibility

Is there adequate commitment from project stakeholders (e.g. experts, users, customers)? (PF1)

We have not yet ascertained the level of support from external stakeholder groups, e.g. military and humanitarian agencies, with regard to the outcomes of this project. We believe that such feedback can be obtained once a prototype system has been developed to adequately demonstrate the capabilities of the system.

Efforts to secure the cooperation of domain experts are ongoing.

Are the required knowledge and competences available? (PF2)

Our research group is at the forefront of efforts to develop, investigate and exploit the semantic web. We also have extensive experience of delivering high-quality knowledge engineering solutions in a variety of application domains including search and rescue (Cottam & Shadbolt, 1998), defence-related decision support (e.g. FOAS, FOAEW), corporate knowledge management (e.g. Scottish Amicable) and biomedicine (e.g. MIAKT – see Shadbolt et al, 2004b). Members of our team are fully conversant with the technical requirements of the current project.

Are the expectations regarding the project and its results realistic? (PF3)

We have tried to be explicit about the core capabilities of the proposed system and the project outcomes in the context of this report (see, for example, Section 3). We suspect the capability of the system to be of approximately equal complexity to other semantic web initiatives undertaken by our group, e.g. CS AKTiveSpace[71] (winner of the 2003 Semantic Web Challenge[72])

10.3 Technical Feasibility

How complex, in terms of knowledge stored and reasoning processes to be carried out, is the task to be performed by the considered knowledge system solution? Are state-of-the-art methods and techniques available and adequate? (TF1)

We believe the knowledge structures required to support the desired knowledge processes are of sufficiently complexity to tax the current-the-state-art in terms of knowledge engineering techniques, without jeopardising the ability of the project to realize its intellectual and technological commitments. We aim to exploit tried and tested techniques for knowledge capture most of which are accessible through robust knowledge acquisition environments such as Protégé and PCPACK. We also intend to rely heavily on the CommonKADS methodology since this provides a structured approach to knowledge engineering based on two decades of research activity and industry use.

[71] http://www.aktors.org/technologies/csaktivespace/

[72] http://challenge.semanticweb.org/

Perhaps the key technical risk here concerns the availability of subject matter experts who can support the knowledge capture process. While every effort has been made to provide comprehensive coverage of the target domain, it is by no means clear whether the captured knowledge is similar to that acquired by experienced operatives in the area of humanitarian aid deployment. Attempts have been made to elicit the voluntary co-operation of a number of aid organisations, but these efforts have, so far, proved unsuccessful. One means of addressing this concern relates to the provision of adequate knowledge editing facilities that will enable the knowledge infrastructure of the application to be updated and refined by those outside the knowledge engineering community. Given this capability we suspect that any issues concerning the integrity of acquired knowledge can be adequately addressed during the course of the project lifecycle.

Are there critical aspects involved, relating to time, quality, needed resources, or otherwise? (TF2)

It is important for the system to comply with key knowledge management objectives such as delivering information at the right time, in the right form, in the right place with the right quality. One issue here concerns the timeliness of information provision. It is not clear, at the moment, what processing demands will be placed on the system and whether these will cause the system to violate temporal constraints surrounding the timeliness of decision support and information provision. Past work has shown that in real-time decision support systems processing speed is often a limiting factor in the operational integrity of the system. This is particularly so when certain types of expert system shell are used, e.g. CLIPS. As such, a requirement for real-time response profiles may impose hard constraints on the choice of implementation technology.

Another issue concerns access to information sources about events of strategic relevance to the effective deployment of humanitarian aid programmes. Some content providers, e.g. news agencies, may require subscription and financial remuneration in return for access to their services. In other cases information access may be restricted due to security constraints and concerns, an issue that also arises, in a related form, in the context of information dissemination and knowledge processing, e.g. should an agent be privy to information that was originally derived (perhaps in part) from restricted information that they would not normally have access to.

Finally, the quality of the information provided by the system is of paramount significance, especially in military domains where the choice between response options may have a number of politically sensitive repercussions. As such we should aim to avoid over-reliance on the system by carefully qualifying all information with respect to certainty criteria. In cases where such certainty cannot be adequately ascertained the system should err on the side of caution.

Is it clear what the success measures are and how to test for validity, quality and satisfactory performance? (TF3)

A number of evaluation metrics are outlined in Section 6 of this report; however, while it is relatively clear what questions should be asked of a system in terms of performance criteria, it is more difficult to specify how a system should be empirically evaluated with respect to these criteria. Key problems relate to the operationalization of notions such as situational awareness, operational effectiveness, the quality of decision outcomes, and so on. Unless such constructs can be adequately operationalized, then assessments of performance quality and added value are somewhat problematic to say the least.

A related problem concerns the issue of experimental evaluation. A number of specific ideas for experiments have been proposed for the current project, each of which aims to assess the impact of the system on performance metrics and user cognition. The results from these experiments would be suitable for dissemination within the HMI, decision support and augmented cognition communities. As such, it is important to identify any risks and uncertainties associated with experimental analyses at an early stage of the project life cycle. Possible experimental manipulations include adjustments to the level of information fusion provided to subjects, the introduction of progressively noisy information (e.g. conflicting information), different workload levels created by scenarios differing in terms of their complexity, the impact of different types of visualization or interface design on task performance, and so on. Of course, a number of the terms used in the description of experimental evaluations are just as esoteric as their evaluation metric counterparts. For example, what exactly is meant by the term 'information fusion level'? It is clearly important to establish some form of stakeholder consensus about such definitions before appropriate experimental evaluations can be made.

How complex is the required interaction with end users (user interfaces)? Are state-the-art methods and techniques available and adequate? (TF4)

It is not clear at the present time what the exact nature of the user interfaces will be. In all likelihood we will aim to provide a highly customisable visualization environment that can be adapted for different subscribed users according to their problem-solving profile and operational role. It may be that some experimental evaluation is required here, i.e. to investigate the impact of different visualizations on situational awareness and task performance. In some cases, of course, a relatively invariant interface may be required. This is particularly so in the case of military personnel who have become accustomed to working with a standard mission system interface or set of iconographic conventions. In some cases, compliance with these interface conventions may even be mandated by commanding authorities.

It is not even clear that the user interfaces to be developed in the context of the current project must necessarily target the visual modality. A number of projects in 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. Although we do not expect to consider multi-modal issues in the context of the current initiative, it is important to reflect on the variety of interfaces that could be used for system interaction and the impact this has on system processing.

How complex is the interaction with other information systems and other resources (interoperability, systems integration)? (TF5)

In general we believe systems-interoperability issues to be adequately addressed by the endorsement of industry standard formats for information exchange, e.g. XML, RDF, SOAP, WSDL, etc. One potent concern is the sourcing of information from content providers, some of which may require financial remuneration for information access and others which only provide access to relatively unstructured information. The latter of these issues can, we believe, be adequately addressed by exploiting the current state-of-the-art in natural language processing and intelligent text analysis. To this end we have established contact with another DIF DTC group specifically concerned with intelligent information extraction from largely unstructured textual sources (DIF DTC Project 7.5).

10.4 Summary

The current section has attempted to highlight a number of risks and uncertainties that impact on the feasibility of the proposed solution. In most cases these issues are adequately addressed within the resource framework of the project and no further project actions are required. Outstanding concerns include:

1. the operationalization of constructs for system evaluation (TF3)
2. the cooperation of subject matter experts (TF1, PF1)
3. customer expectations regarding system functionality (PF1)
4. the appropriate evaluation of temporal constraints (TF2)
5. the need to adequately address security concerns relating to information provision (TF2)

In response to these concerns we recommend the following project actions:

1. arrange stakeholder meetings to establish consensus on the operationalization of evaluation metrics and the meaning of analytical terms
2. continued programme of effort to contact humanitarian aid agencies and enlist support where possible
3. contact interested stakeholder groups (e.g. military and humanitarian aid agencies) with a view to providing early demonstrations of system functionality – elicit feedback and evaluation in the context of such meetings
4. evaluate the violation of temporal constraints in the context of the system prototype – implement any technology changes as required
5. assume that information inferred or derived from restricted information should not be presented to unauthorized subscribers

We fully expect the aforementioned concerns to be adequately addressed by the implementation of these actions.

Appendix A Acronyms and Abbreviations

AAAI	American Association for Artificial Intelligence
AEW	Airborne Early Warning
AI	Artificial Intelligence
AKT	Advanced Knowledge Technologies
API	Application Programmatic Interface
ASP	Active Server Pages
BT	British Telecommunications
C4ISTAR	Command, Control, Communications, Computers, Information/Intelligence, Surveillan
CAP	Combat Air Patrol
CFS	Combat Flight Simulator
CLIPS	C Language Integrated Production System
CML	Conceptual Modelling Language
CommonKADS	Common Knowledge Analysis and Design System
COTS	Commercial Off-The-Shelf
CSS	Cascading Style Sheet
DAML	DARPA Agent Markup Language
DARPA	Defence Advanced Research Projects Agency
DIF	Data and Information Fusion
DL	Description Logic
DTC	Defence Technology Centre
ECAI	European Conference on Artificial Intelligence
EKAW	European Knowledge Acquisition Workshop
ERS	Electromagnetic Remote Sensing
ESWS	European Semantic Web Conference
FLIR	Forward-Looking Infrared Radar
FOAEW	Future Organic Airborne Early Warning
FOAS	Future Offensive Air System
GSS	Graph Stylesheets

HFI	Human Factors Integration
HMI	Human Machine Interface
HMM	Hidden Markov Models
HP	Hewlett-Packard
HP	Hewlett-Packard
HTML	Hyper Text Markup Language
IAM	Intelligence, Agents and Multimedia
ICT	Information and Communication Technology
ICWS	International Conference on Web Services
IEEE	Institute of Electrical & Electronics Engineers
IJCAI	International Joint Conference on Artificial Intelligence
IKAT	IKEW Knowledge Acquisition Toolkit
IKEW	Internet-enabled Knowledge Elicitation Workbench
ILP	Inductive Logic Programming
IR	Infra-Red
IRST	Infra-Red Search and Track
ISSCTMC	Infinite State Space Continuous Time Markov Chains
ISTAR	Information/Intelligence, Surveillance, Targeting Acquisition and Reconnaissance
ISWC	International Semantic Web Conference
JDBC	Java Database Connectivity
JESS	Java Expert System Shell
JSP	Java Server Pages
KA	Knowledge Acquisition
KBS	Knowledge-Based System
MANDRIL	Message Analysis and Data Reduction for the Integration of Links
MIAKT	Medical Imaging and Advanced Knowledge Technologies
ML	Machine Learning
MOD	Ministry of Defence
MOE	Measures of Effectiveness
MOKA	Methodology for Knowledge-Based Engineering Applications
MOP	Measures of Performance
MRP	Markov Regenerative Processes
NA	Northern Alliance

NEC	Network Enabled Capability
NLP	Natural Language Processing
OIL	Ontology Inference Layer
OWL	Web Ontology Language (W3C)
PCPACK	PC-Enabled Portable ACquisition of Knowledge
PDA	Personal Digital Assistant
PhD	Philosophiae Doctor
PHP	PHP: Hypertext Preprocessor
PPI	Plane Position Indicator
PSM	Problem-Solving Method
QPR	Quarterly Progress Report
RDBMS	Relational Database Management System
RDF	Resource Description Framework
RDFS	RDF Vocabulary Description Language
RDQL	RDF Data Query Language
RF	Radio Frequency
RuleML	Rule Markup Language
SME	Subject Matter Expert
SOAP	Simple Object Access Protocol
SQL	Structured Query Language
SVG	Scalable Vector Graphics
SWRL	Semantic Web Rule Language
TDS	Technical Demonstrator System
UK	United Kingdom
UML	Unified Modelling Language
URI	Uniform Resource Identifier
USA	United States of America
VBA	Visual BASIC for Applications
W3C	World Wide Web Consortium

WP	Work Package
WSDL	Web Services Description Language
WWW	World Wide Web
XML	eXtensible Markup Language

Appendix B References

1. Beach, K. (1988). The role of external mnemonic symbols in acquiring an occupation. In *Practical Aspects of Memory*. M. M. Gruneberg and R. N. Sykes. New York, Wiley: 342-346.
2. Clark, A. (1997) *Being There: Putting Brain, Body and World Together Again*. Cambridge, Massachusetts, MIT Press.
3. Cottam, H. and Shadbolt, N.R. (1998) Knowledge acquisition for search and rescue planning. *International Journal of Human-Computer Studies*, 48(4): 449-473
4. Gibbins, N.; Harris, S. and Shadbolt, N. R. (2003) Agent-based Semantic Web Services. In *Proceedings The Twelfth International World Wide Web Conference (WWW2003)*, Budapest, Hungary.
5. Harris, S and Gibbins, N. (2003) 3store: Efficient Bulk RDF Storage. In *Proceedings of the First International Workshop on Practical and Scalable Semantic Web Systems (PSSS2003)*, Sanibel Island, Florida, USA.
6. Jackson, P (1999) *Introduction to Expert Systems* (3rd Edition), Addison Wesley Longman Limited, Harlow, Essex, UK
7. Leonard, T. and Glaser, H. (2001) Large scale acquisition and maintenance from the web without source access. In *Proceedings of Workshop 4, Knowledge Markup and Semantic Annotation (K-CAP 2001)*, ACM Press, pages 97-101.
8. Newell, A. (1982) The knowledge level. *Artificial Intelligence*, 18: 87-127
9. Schreiber G., Akkermans, H., Anjewierden, A., de Hoog, R., Shadbolt, N.R., Van de Velde, W. and Wielinga, B. (2000) *Knowledge Engineering and Management: The CommonKADS Methodology*. MIT Press, Massachusetts
10. Shadbolt, N. R.; Gibbins, N.; Glaser, H.; Harris, S. and schraefel, m. (2004a) CS AKTive Space: or how we learned to stop worrying and love the Semantic Web. IEEE Intelligent Systems, 19(3): 41-47
11. Shadbolt, N. R., Lewis, P., Dasmahaptra, S., Dupplaw, D., Hu, B. and Lewis, H (2004b). MIAKT: combining grid and web services for collaborative medical decision making. In *Proceedings of the 2nd e-Science All Hands Meeting*, Nottingham, England, UK.
12. Shadbolt, N.R., O'Hara, K., Crow, L. (1999) The Experimental Evaluation of Knowledge Acquisition Techniques and Methods: History, Problems and New Directions. *International Journal of Human-Computer Studies*, 51(4): 729-755
13. Stefik, M (1995) *Introduction to Knowledge Systems*, Morgan Kaufmann, San Francisco, California, USA