Rule-Based Intelligence on the Semantic Web

Implications for Military Capabilities

Dr Paul Smart

Senior Research Fellow
School of Electronics and
Computer Science
University of Southampton
Southampton
SO17 1BJ
United Kingdom

26th November 2007

Report Documentation Page

Report Title:	Rule-Based Intelligence on the Semantic Web		
Report Subtitle	Implications for Military Capabilities		
Project Title:	N/A		
Number of Pages:	57	**Version:**	1.2
Date of Issue:	26/11/2007	**Due Date:**	22/11/2007
Performance Indicator:	EZ~01~01~17	**Number of References:**	95
Reference Number:	DT/Report/RuleIntel		
Report Availability:	APPROVED FOR PUBLIC RELEASE; LIMITED DISTRIBUTION		
Abstract Availability:	APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED		
Authors:	Paul Smart		
Keywords:	semantic web, ontologies, reasoning, decision support, rule languages, military		

Primary Author Details:

Dr Paul Smart
Senior Research Fellow
School of Electronics and Computer Science
University of Southampton
Southampton, UK
SO17 1BJ
tel: +44 (0)23 8059 6669
fax: +44 (0)23 8059 2783
email: ps02v@ecs.soton.ac.uk

Client Details:

Abstract:

Rules are a key element of the Semantic Web vision, promising to provide a foundation for reasoning capabilities that underpin the intelligent manipulation and exploitation of information content. Although ontologies provide the basis for some forms of reasoning, it is unlikely that ontologies, by themselves, will support the range of knowledge-based services that are likely to be required on the Semantic Web. As such, it is important to consider the contribution that rule-based systems can make to the realization of advanced machine intelligence on the Semantic Web. This report aims to review the current state-of-the-art with respect to semantic rule-based technologies. It provides an overview of the rules, rule languages and rule engines that are currently available to support ontology-based reasoning, and it discusses some of the limitations of these technologies in terms of their inability to cope with uncertain or imprecise data and their poor performance in some reasoning contexts. This report also describes the contribution of reasoning systems to military capabilities, and suggests that current technological shortcomings pose a significant barrier to the widespread adoption of reasoning systems within the defence community. Some solutions to these shortcomings are presented and a timescale for technology adoption within the military domain is proposed. It is suggested that application areas such as semantic integration, semantic interoperability, data fusion and situation awareness provide the best opportunities for technology adoption within the 2015 timeframe. Other capabilities, such as decision support and the emulation of human-style reasoning capabilities are seen to depend on the resolution of significant challenges that may hinder attempts at technology adoption and exploitation within the 2020 timeframe.

Classification of Report:	Classification of Abstract:	Classification of this Page:
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED

Limited Distribution Notice

Acknowledgements

This work was supported by a contract awarded to the University of Southampton by Detica, London, UK.

Abstract

Rules are a key element of the Semantic Web vision, promising to provide a foundation for reasoning capabilities that underpin the intelligent manipulation and exploitation of information content. Although ontologies provide the basis for some forms of reasoning, it is unlikely that ontologies, by themselves, will support the range of knowledge-based services that are likely to be required on the Semantic Web. As such, it is important to consider the contribution that rule-based systems can make to the realization of advanced machine intelligence on the Semantic Web. This report aims to review the current state-of-the-art with respect to semantic rule-based technologies. It provides an overview of the rules, rule languages and rule engines that are currently available to support ontology-based reasoning, and it discusses some of the limitations of these technologies in terms of their inability to cope with uncertain or imprecise data and their poor performance in some reasoning contexts. This report also describes the contribution of reasoning systems to military capabilities, and suggests that current technological shortcomings pose a significant barrier to the widespread adoption of reasoning systems within the defence community. Some solutions to these shortcomings are presented and a timescale for technology adoption within the military domain is proposed. It is suggested that application areas such as semantic integration, semantic interoperability, data fusion and situation awareness provide the best opportunities for technology adoption within the 2015 timeframe. Other capabilities, such as decision support and the emulation of human-style reasoning capabilities are seen to depend on the resolution of significant challenges that may hinder attempts at technology adoption and exploitation within the 2020 timeframe.

Contents

Figures

1 Introduction

1.1 Visions of the Semantic Web

The Semantic Web (Berners-Lee et al., 2001) is a vision of the future potential of the World Wide Web (WWW) to provide a global infrastructure for the representation, dissemination and exploitation of human knowledge. In some ways the use of the phrase 'Semantic Web' is unfortunate: we already have a global repository of meaningful information that covers practically every area of human endeavour and experience - it is the conventional WWW. Why would we want to go beyond this existing capability? The answer to this question lies in the fact that the WWW is designed primarily for human consumption – most of the information available on the Web is meaningful only in the sense that humans are able to interpret it. What makes the Semantic Web different from the conventional Web is the emphasis it places on representational formalisms that make the meaning of information content explicit. Once we have a framework for unambiguously representing the meaning of information, then we have a framework within which intelligent systems are able to manipulate and exchange information content in a semantically-coherent and semantically-sensible fashion. The kinds of capabilities that may be supported by the Semantic Web are still the subject of considerable speculation; however, the following capabilities seem to be at least theoretically plausible:

1. Improved search and retrieval capabilities grounded in the fact that humans can better communicate their interests and intent to machines (the Semantic Web will enable agents to explicitly specify what they want to find in particular information retrieval contexts, something which goes beyond the capability engendered by plain keyword searches).
2. Improved inter-operability between disparate systems, especially in relation to information exchange, knowledge transfer and collaborative problem-solving.
3. Improved aggregation of information content either at the physical or virtual level, i.e. an ability to aggregate distributed information content for the purposes of specialized knowledge portals and services.
4. Improved clustering and organization of information content with respect to dynamically specified categories of interest, e.g. an ability to dynamically reorganize document repositories in ways that reflect the idiosyncratic interests and perspectives of end-user agents.
5. Improved knowledge discovery and creation, including an ability to use a combination of data mining techniques, statistical analysis and reasoning to discover new contingencies and statistical dependencies in large datasets (consider the transformative potential of an ability to publish scientific data on the Web and make that data amenable to automated knowledge processors[1]).

This is by no means an exhaustive list, but it does serve to convey some of the flavour of the Semantic Web vision.

The first step in the development of the Semantic Web is the availability of a knowledge representation language that can be used to express human knowledge in a form that is amenable

[1] See also Berners-Lee & Hendler (2001).

to machine processing. Recently, efforts to provide such a language have coalesced around the Web Ontology Language (OWL) (Antoniou & van Harmelen, 2004). OWL provides a language for describing the semantic infrastructure of a domain of discourse. As can be seen from Figure 1, which depicts the "Semantic Web Layer Cake" (a popular representation of the architectural components of the Semantic Web), the ontology representation language is built-on top of the Resource Description Framework (RDF) and RDF Schema. Above the ontology layer are the logic and proof layers, each of which provides a foundation for rule-based processing. It is these components of the Semantic Web (and the reasoning capabilities they support) that provides the primary focus of analysis for the current report.

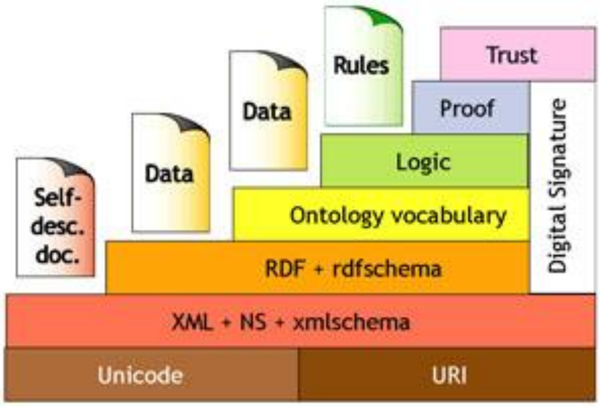

Figure 1: The Semantic Web Layer Cake

1.2 Reasoning on the Semantic Web

Since the first publication of the "Semantic Web Layer Cake", the role of rules in the Semantic Web has been somewhat controversial. Part of the reason for this controversy concerns the prevailing sense of confusion about the kinds of capabilities that rules are intended to provide, or indeed what types of knowledge they are designed to capture. As Allemang (2006) notes:

> *"Rules have sometimes been given a central role, at other times a peripheral role, and sometimes left out completely. Why such variation for a technology with thirty years of background? The reason for these differences of opinion stem from different goals for the inclusion of rules in the Semantic Web stack. At one extreme are the Description Logicians who see no need for a general-purpose programming language in the Semantic Web stack. At the other extreme are those who want to build a web infrastructure with the capacity for emergent intelligence." (Allemang, 2006)*

In some cases, the confusion about the role and impact of rules is clearly justified. It is, as yet, unclear what kinds of capabilities will be supported by the deployment of reasoning capabilities within the Semantic Web. Furthermore, it would be unfair to expect a coherent vision of rule-based capabilities in the absence of some agreement about the technological underpinnings of such a capability. The fact is that the development of rules and reasoning capabilities for the Semantic Web is still in its infancy. While some rule languages for the Semantic Web have been developed, e.g. Rule Markup Language (RuleML) and Semantic Web Rules Language (SWRL), there is little consensus at the present time regarding how rules should be represented in the Semantic Web, or indeed the kinds of capabilities they should support.

The knowledge representation languages of the Semantic Web, RDF and OWL, are based on a subset of predicate logic, called first-order logic. The formal semantics of OWL support a particular type of reasoning, called subsumption reasoning. This essentially allows a system to infer that one class of objects is a subset of (is subsumed by) another set based on the logical characterizations made of the concepts in question. The following is an example of a rule that captures an inference supported by the semantics of the ontology representation language:

```
(defrule transitivity (declare (salience 100))
   (triple (predicate ?p) (subject ?x) (object ?y))
   (triple (predicate ?p) (subject ?z) (object ?x))
   (triple
           (predicate "http://www.w3.org/1999/02/22-rdf-syntax-ns#type")
           (subject   ?p)
           (object    "http://www.w3.org/2002/07/owl#TransitiveProperty")
   )
   =>
   (assert (triple (predicate ?p) (subject ?z) (object ?y)))
   )
```

Figure 2: Transitivity Inference Rule

This rule exploits the semantics of transitive properties to infer that if *x* is related to *y*, and *z* is related to *x*, and the property that links these objects is a type of transitive property, then it must be the case that *z* is related to *y* by the same property.

Reasoners that provide support for Description Logic (DL) reasoning[2] include Pellet (Parsia & Sirin, 2004; Sirin et al., 2007) and Racer (Haarslev & Möller, 2003). These reasoners are typically integrated into knowledge editing environments, such as Protégé (Noy et al., 2001), and they assist with the ontology development process by performing logical consistency checks and semantic validation services. Unfortunately, however, the kind of reasoning these reasoners can perform is limited to the semantics of the knowledge representation language; they cannot be used to compute certain types of relationships. For example, the computation of a *hasUncle* relationship requires an ability to conditionally assert a new relationship based on the existence of at least two other relationships: *hasParent* and *hasBrother*. In general, description logics are unable to express chains of joins across different predicates (Antoniou & Wagner, 2003), and this limits the kind of inferences that they can participate in.

To complement the reasoning capabilities supported by DL reasoners, a number of attempts have been made to extend the expressivity of OWL with Horn logic rules. SWRL (Horrocks et al., 2004) is one of the best examples of this approach. It allows users to write Horn-like rules that can be expressed in terms of OWL concepts and that can then be used to reason about OWL individuals.

In some cases, a reasoning process that initially looks as though it might require the services of Horn-logic rules can be delegated to a DL reasoner by maximally exploiting the semantic axioms of the ontology representation language. As an example of this capability, suppose we want to use an ontology to provide decision support in respect of the relative priorities of mine hazard areas for

[2] In the context of this report, I will refer to any form of reasoning that is supported by the semantics of the ontology representation language as DL reasoning.

demining operations[3]. We can delegate this task to a DL reasoner by creating a new class within the ontology to represent our reasoning goal (in this case the new class is called HighPriorityMineHazardArea). We can then use the logical expressions provided by the ontology representation language to specify what we *mean* by this concept. The actual definition of the class is depicted in Figure 3[4]. Now that we have communicated what we mean by a high priority mine hazard area, in a form a machine can understand, we can ask a DL reasoner to classify all the mine hazard area instances that fulfil the membership criteria of the target class. This will return a list of mine hazard areas that are of high priority for humanitarian demining.

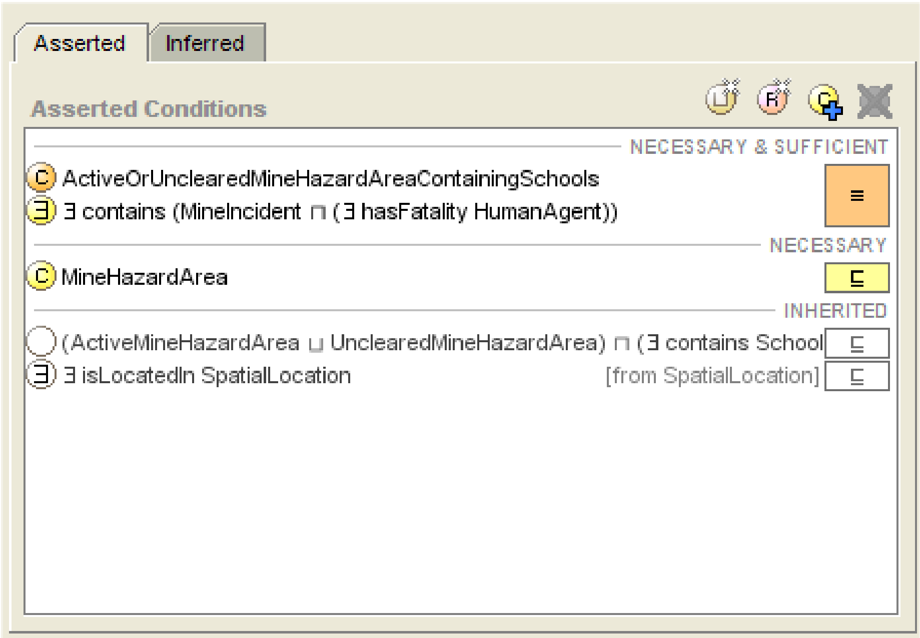

Figure 3: HighPriorityMineHazardArea Class

1.3 The Military Perspective

The military environment presents a number of challenges for the Semantic Web. Specific challenges include the need to deal with the idiosyncrasies of the military network environment (e.g. its mobile, *ad hoc* and wireless nature), the need to engender effective and robust information exchange solutions via ontology alignment/mapping mechanisms, the need to cope effectively with uncertain, incomplete and unreliable information, and the need to deal with the threat posed by malign agencies that may exploit the open nature of the Semantic Web to subvert coalition decision-making processes (see Smart & Shadbolt, 2007). Another notable challenge, and one that is the central focus of this report, is the need to develop robust reasoning capabilities that are able to assist military decision-makers with respect to a variety of tasks, such as situation assessment, sense-making and mission planning. Although the vision of the Semantic Web, and the reasoning capabilities it will eventually support, are largely commensurate with the requirements of military agencies vis-à-vis future defence capabilities, a number of critical challenges still need to be addressed. These include,

[3] This example is taken from the Data and Information Fusion Defence Technology Centre (DIF DTC) SEMIOTIKS initiative. SEMIOTIKS aims to provide semantically-enabled capabilities in support of humanitarian demining and Explosive Ordnance Disposal (EOD) operations.
[4] A high priority mine hazard area is defined in this case as an uncleared mine hazard area that contains at least one school and which is also is the location of at least one mine-related fatality.

above all, the need to provide efficient and scalable reasoning technologies that deliver decision outcomes within an operationally-relevant timeframe. This is a significant challenge because extant reasoning systems often show poor scalability and performance with respect to knowledge-intensive tasks, especially synthetic task types, such as planning, design, prediction, simulation, etc. (see Schreiber et al., 2000). Concerns about performance and scalability will no doubt be magnified by the need to operate within military network environments. The time required to dynamically aggregate relevant information across the multiple nodes of a resource constrained network, for example, may contribute to a performance impairment that exceeds the acceptability threshold for time-critical missions.

This report aims to review the current state-of-the-art with respect to semantic reasoning systems and the capabilities they support. Section 2 provides an overview of the various technologies currently available to support reasoning in the context of the Semantic Web. This section summarizes the current state-of-the-art with respect to rule languages and rule engines. Section 3 presents a number of military application areas where reasoning capabilities would be of potential relevance. These areas (e.g. decision support) are pretty much those that have been explored in the context of conventional expert systems research, although some capabilities, e.g. semantic integration and interoperability, are pretty much unique to the Semantic Web. The barriers to technology adoption within a military context are described in Section 4. The aim here is to explicate the factors that may limit the perceived utility or acceptability of semantic reasoning solutions within the military community. Section 5 presents a tentative timeframe for the adoption of semantic reasoning technologies with respect to a number of application areas. This analysis is based on a number of factors, including the level of maturity of underpinning technologies and the likelihood and early resolution of outstanding technical difficulties. Finally, Section 6 presents some general conclusions and discusses the role played by the defence domain as a proving ground for semantic technologies.

2 Rule-Based Technologies

Reasoning systems are comprised of a range of technological components, including rules, rule languages and rule engines. This section provides a brief overview of these technological components, particularly those that have been studied in the context of the Semantic Web.

2.1 Rules

A rule is the basis of inference execution in a knowledge-based system, at least one that operates over semantically-transparent (see Clark, 1989) symbolic atoms. Within the Semantic Web, a rule typically represents a logical entailment between a set of formulas called premises and an assertion called a conclusion. The general form of a rule is as follows:

$$A_1, \ldots A_n \rightarrow B$$

where Ai and B are atomic formulas.

If a set of rules can be used to infer ANY valid conclusion from the knowledge base, then the inference solution provided by the rules is regarded as complete; if the rule set never permits the assertion of invalid conclusions, then the inference solution is additionally regarded as sound.

Knowledge representation systems based on predicate logic and its various specializations, such as OWL and RDF, are committed to monotonicity assumptions. This means that a DL reasoner can never make inferences that are invalidated by the assertion of additional information – if we know that *x* is an instance of *A*, then the assertion of more information about *x* can never cause it NOT to be an instance of *A*. As such, DL reasoners rely on monotonic rules to compute the entailments implied by the logical axioms of the ontology language – all DL reasoners are, in essence, monotonic reasoning systems.

Family	Gills	Spore Colour	Cap Shape	Stem Width	Poisonous
Stropharia	Free	White	Umbrella	Broad	Poisonous
Russula	Free	White	Coral	Narrow	Not-Poisonous
Amanita	Adnate	White	Disk	Broad	Poisonous
Amanita	Adnate	Black	Umbrella	Broad	Poisonous
Bolete	Adnate	White	Umbrella	Broad	Poisonous
Ink Cap	Sinuate	Pink	Umbrella	Broad	Not-Poisonous
Bolete	Decurrent	Brown	Finger	Narrow	Not-Poisonous
Amanita	Adnate	Pink	Umbrella	Broad	Not-Poisonous

Figure 4: Sample Dataset for Mushroom Classification Service

Another possible interpretation of rules is that they are representational formalisms that essentially capture the predictive contingencies or statistical regularities within a domain of discourse. According to this view, rules represent a special type of relationship in which an association (largely probabilistic in nature) is established between one or more properties of domain concepts. Take for example, the data presented in Figure 4, which describes the characteristics of various mushroom families. Using this data, we can use a rule induction engine to compute the following rule:

```
IF Mushroom.Gills = 'Adnate' THEN Mushroom.Poisonous = 'Poisonous' (0.75)
```

The value of 0.75 in this case indicates the probability that the consequent (the THEN part of the rule) is correct. This rule can be interpreted in the following terms: “if a mushroom has adnate[5] gills, then there is a 75% chance that it is also poisonous”. The knowledge captured by this rule statement is typical of real-world knowledge and it underpins our problem-solving competency in a whole variety of naturalistic decision-making contexts. Unfortunately, many of the rule systems currently seen within the Semantic Web cannot represent probabilistic knowledge of the type captured by this rule. Instead they are largely committed to inferring valid (i.e. true) conclusions from premises that are similarly valid. This is something of a weakness in terms of the ability of semantic rule languages to adequately cope with the vagaries of real-world knowledge (see Section 4.2).

Following the analysis of rule languages in Boley et al (2007), we identify three categories of rules: deductive rules, reactive rules and normative rules. A brief overview of these rule types is presented in subsequent sections.

2.1.1 Deductive Rules

Deductive rules allow a reasoning system to infer new knowledge on the basis of existing knowledge. Deductive rules are sometimes referred to as constructive rules because they are able to exploit knowledge-rich contingencies to create new knowledge; they thereby enrich the epistemic substrate for further reasoning processes (Bry & Marchiori, 2005)[6].

Deductive rules are often comprised of two parts: a ‘head’ and a ‘body’. The ’head’ part typically provides a specification of the data to be constructed or inferred, and the ‘body’ part queries the underlying knowledge bases(s). Both parts of the rule usually share variables (commonly prefixed with a question mark, e.g. ‘?entity’. The variables get bound to data items in the ‘body’ and these bindings are then used in the ‘head’ to assert new data (see Figure 5[7]).

[5] The term ‘adnate’ refers to a particular morphological associated of mushroom gills with the mushroom stem.

[6] Deduction rules are also called derivation rules in the business rules community, constructive rules by logicians, and views in the database community (Boley et al., 2007).

[7] This rule is a CLIPS rule taken from the Future Offensive Air System (FOAS) pilot aiding domain (see Shadbolt et al, 2000).

```
(defrule generate-appropriate-threat-reaction-against-threat-more-than-range
  (active [threat-reaction-rule])
  (object (radar-mode NA)
          (missile-mode NA)
          (min-range ?min-range&~none)
          (max-range none)
          (is-a threat-reaction)
          (name ?reaction)
          (action-type ?action-type)
          (threat-type ?threat-type))
  (object (slot type)
          (value ?threat-type)
          (instance ?threat)
          (is-a single-value-datum)
          (name ?threat-type-datum))
  (object (name ?threat)
          (entity ?entity))
  (object (instance ?entity)
          (slot tactical-range)
          (maximum ?max&:(> ?max ?min-range))
          (is-a maximum-datum)
          (cf ?max-cf)
          (name ?max-range-datum))
  (not (object (instance ?entity)
               (slot tactical-range)
               (cf ?val-cf&:(> ?val-cf (/ ?max-cf 2)))
               (is-a single-value-datum)))
=>
  (bind ?basis (create$ ?threat-type-datum
                        ?max-range-datum))
  (if (> (lowest-cf ?basis) 20) then
  (bind ?reason (str-cat (send ?action-type get-name-string) " is effective"
                  " against " ?threat-type
                  " at more than " ?min-range "."))
  (bind ?rationale (create-KBS-source ?reason 25 $?basis))
  (make-instance (gensym)
      of appropriate-threat-reaction
    (threat-reaction ?reaction)
    (threat ?threat))))
```

Figure 5: Example of Variable Binding in a Deductive Rule

The semantic query language, SPARQL[8], provides a CONSTRUCT clause which can be used to create new RDF triples from RDF datasets (see Figure 6[9]). This type of SPARQL query can be viewed as a deductive rule (Polleres, 2007) because the query is being used to derive new knowledge from previously asserted facts[10]. Other examples of rule languages based on deductive rules are SQL views, Datalog, Prolog, and most other logical rule languages (Boley et al., 2007). XSLT templates can also be viewed as deductive rules, as can queries that follow the XQuery language (Bry & Marchiori, 2005).

```
CONSTRUCT
{
        ?mine sem:hasType sem:APERS_PI-MI-SR .
        ?mine sem:hasTypeP “.50”^^xsd:float .
}
WHERE
{
        ?incident rdf:type sem:MineIncident .
        ?incident sem:isLocatedIn sem:Afghanistan .
        ?incident sem:involves ?mine .
        ?mine rdf:type sem:Mine .
        ?mine sem:usesExplosive sem:TNT.
}
```

Figure 6: SPARQL CONSTRUCT Query

[8] http://www.w3.org/TR/rdf-sparql-query/

[9] This query is taken from the domain of explosive ordnance disposal.

[10] This view of SPARQL queries as rules suggests a potential strategy for improved rule processing efficiency. If rules can be recast as queries, then we can delegate rule processing to query engines that, in many cases, are more optimized than their rule engine counterparts.

2.1.2 Reactive Rules

Reactive rules are the basis of reactive systems that can respond to the occurrence of specific events. Two types of reactive rule are generally recognized: Event Condition Action (ECA) rules and production rules. ECA rules (Papamarkos et al., 2003) take the form of ON Event IF Condition DO Action, which specifies that the Action should be executed automatically when the Event is detected, (providing the Condition holds, of course). Another category of reactive rules are the production rules. These rules form the basis of production rule systems such as Jess[11], CLIPS[12], JRules[13], JBoss Rules[14] and OPS5[15]. In this case, the rule takes the form of WHENEVER Condition DO Action. During the course of inference execution, the conditional statements of a production rule are continuously evaluated, and selected data is then used to execute the actions specified in the Action part of the rule.

Reactive rules are seen as a means of enabling the transition from a largely passive Web, where data sources can only be accessed to obtain information, and a (future) dynamic Web, where data sources are enriched with reactive behaviour (Berstel et al., 2007).

2.1.3 Normative Rules

Normative rules are rules that serve to constrain the data values or logic of an application. The classical example is that of an integrity constraint in traditional relational database systems, e.g. each customer must have a unique name. Data schemas, especially tree grammars in their various disguises, e.g. DTD, XML Schema, RelaxNG, etc., express normative rules (Bry & Marchiori, 2005).

2.2 Rule Languages

There are many rule languages in use today, each with its own syntactic and semantic idiosyncrasies. In this section we provide an overview of the rule languages that have been most extensively studied in the context of the Semantic Web.

2.2.1 RuleML

RuleML (Boley et al., 2001) is a markup language for the Semantic Web that was proposed by the Rule Markup Initiative[16]. RuleML provides an XML language for rule representation that covers all the rule types described in Section 2.1.

2.2.2 SWRL

Like RuleML, SWRL[17] is a candidate rules language for the Semantic Web. It builds on both OWL (specifically, the OWL Lite and OWL DL sub-languages) and the unary/binary datalog sub-languages of RuleML (O'Connor et al., 2005). The SWRL proposal extends the set of OWL axioms to include Horn-like rules, and it thus enables Horn-like rules to be combined with an OWL knowledge base. SWRL rules take the form of an implication between an antecedent (body) and consequent (head). The intended meaning of the rule can be read as: whenever the conditions specified in the antecedent hold, then the conditions specified in the consequent must also hold. Both the

---

[11] http://herzberg.ca.sandia.gov/
[12] http://www.ghg.net/clips/CLIPS.html
[13] http://www.ilog.com/products/jrules/
[14] http://www.jboss.com/products/rules
[15] http://en.wikipedia.org/wiki/OPS5
[16] http://www.ruleml.org/
[17] http://www.w3.org/Submission/SWRL/

antecedent (body) and consequent (head) may consist of zero or more atoms. Multiple atoms in either the head or body of a rule are treated as a conjunction of the rule atoms, although rules with conjunctive consequents can always be transformed into multiple rules, each with an atomic consequent.

As previously mentioned, SWRL is built on top of OWL, which allows SWRL rules to be represented as OWL individuals. This is important because it allows SWRL to exploit any technology developed to support the processing of RDF/OWL ontologies. Examples include editors, parsers, knowledge bases, query engines and other rule systems. Like OWL, SWRL has an XML presentation syntax that combines elements from both the ruleml and swrlx namespaces (see Figure 7)

```
<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 7: SWRL XML Syntax

SWRL is an attractive language for the Semantic Web. Its syntactic simplicity encourages technology adoption by reducing the training curve for technology incumbents. Users are also supported in the creation and editing of SWRL rules by the availability of rule editors such as the Protégé SWRL Editor (O'Connor et al., 2005) (see Figure 8). This editor is an extension to the popular Protégé-OWL Ontology Editor (Holger et al., 2004) and it enables users to create rules with respect to any domain ontology that is loaded in the Protégé-OWL environment. When editing rules, users can refer directly to the OWL classes, properties and individuals within the loaded ontology and this, in combination with the syntax checking and graphical editing features, greatly simplifies the process of building a semantic rule base.

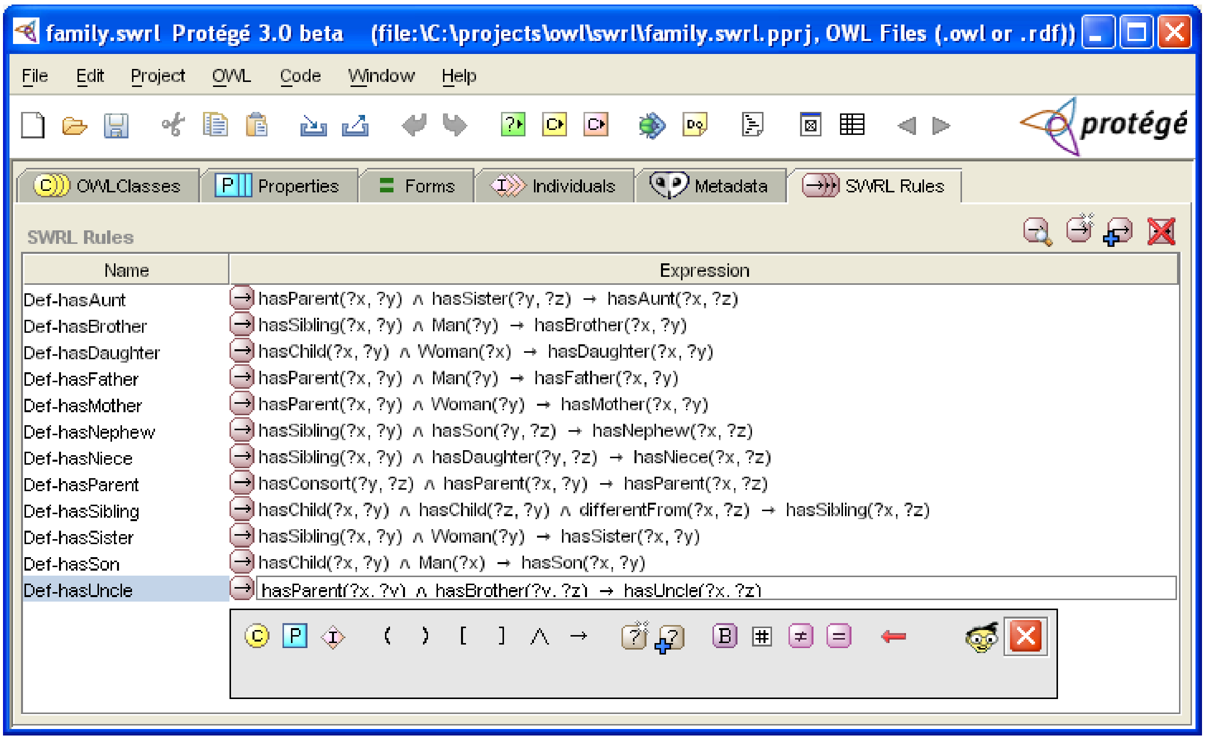

Figure 8: Protégé SWRL Editor

A built-in function library for SWRL allows it to support a variety of function predicates. SWRL rules can also be used to query OWL ontologies using an extended library of SQL-influenced built-in functions (O'Connor et al., 2007). For example, the following rule retrieves all persons in an ontology whose age is less than 5:

```
Person(?p) ^ hasAge(?p,?a) ^ swrlb:lessThan(?a,5) → query:select(?p,?a)
```

The SWRL query library provides basic counting, aggregation, ordering, and duplicate elimination operators.

The SWRL specification does not impose any restriction on how a reasoning process is to be performed using the rules. This means that any number of rule engines could be used to implement the reasoning process, although at present most implementations seem to rely on the Jess rule engine (Golbreich & Imai, 2004; O'Connor et al., 2005) (see Section 2.3.3).

2.2.3 RIF

Information sharing and inter-operability are key themes of the Semantic Web. The emergence of a distributed network of knowledge that is semantically transparent to both humans and machines heralds the promise of an age in which knowledge processes may be increasingly delegated to automated systems (Berners-Lee et al., 2001). The realization of this Semantic Web vision is, however, largely dependent on the development of adequate knowledge representation formalisms (as well as the transfer of human knowledge into a Web-based environment using those formalisms). Rules are an essential ingredient of this vision. They provide much of the inferential power that makes knowledge processing possible, and, as such, it is important to consider the ways in which rules are represented on the Semantic Web. Representational issues are important for two reasons: firstly, they dictate the expressiveness of the rule, which is linked to its inferential potency; secondly, they dictate the extent to which rules can be shared between disparate agent communities, i.e. they dictate the inter-operability of rule processing systems. Interoperability concerns are of paramount importance here because they enable rules to be exploited across the entire knowledge infrastructure of the Semantic Web. This contributes to the progressive enrichment of knowledge systems throughout the Semantic Web by virtue of their ability to engage in rule exchange and merging operations.

Realizing the importance of rule interoperability, the W3C launched the Rule Interchange Format (RIF) Working Group[18] at the end of 2005. This group has a charter to standardize a common format for rule interchange on the Web, but as Boley et al (2007) point out, there are a number of challenges facing the working group. Particular problems relate to the heterogeneity of existing rule languages and the conceptual incompatibilities between existing Semantic Web standards, such as OWL and RDF.

2.3 Rule Engines

This section describes a limited subset of rule engines that have been used to perform semantic reasoning, i.e. reasoning with RDF or OWL ontologies. The most popular reasoners for DL reasoning are Pellet[19], Racer[20] and FaCT++[21]; however, DL reasoning capabilities have also been implemented and tested in a variety of other rule engines, most notably the CLIPS and Jess systems. Although these latter systems are not optimized for DL reasoning, their flexibility with respect to rule-based processing makes them an attractive alternative for implementing semantic reasoning capabilities.

2.3.1 Pellet, Racer and FaCT++

Pellet (Parsia & Sirin, 2004; Sirin et al., 2007), Racer (Haarslev & Möller, 2003) and Fact++ (Tsarkov & Horrocks, 2006) are well known DL reasoners that are commonly used to support ontology development in conjunction with the Protégé knowledge editing environment. All three systems implement tableaux algorithms (Baader & Sattler, 2001), and thus capitalize on the extensive research that has been undertaken within the description logics community. The three systems have more or less similar capabilities and limitations, although FaCT and FaCT++ do not directly support ABox reasoning[22].

2.3.2 CLIPS

CLIPS[23] is an expert system shell that combines a rule-based inference engine with object-oriented and procedural programming facilities (Giarratano & Riley, 1994). These features make CLIPS a highly versatile environment for rule-based programming, and this is borne out by its successful application in a number of problem domains, including mission planning systems for Airborne Early Warning (AEW) (Smart, 2002) and decision support systems in military fighter aircraft (Shadbolt et al., 2000). For the most part, CLIPS has been used to implement classical expert systems, which is not necessarily surprising given that its development predates the Semantic Web[24] (and to some extent

[18] http://www.w3.org/2005/rules/
[19] http://pellet.owldl.com/
[20] http://www.racer-systems.com/
[21] http://owl.man.ac.uk/factplusplus/
[22] A DL knowledge base consists of two components: a TBox and an ABox. The TBox describes the terminology, while the ABox contains assertions about individuals. Correspondingly, DL reasoning includes TBox reasoning (i.e., reasoning with concepts) and ABox reasoning (i.e., reasoning with individuals).
[23] http://www.ghg.net/clips/CLIPS.html
[24] CLIPS was created in 1985 - http://www.ghg.net/clips/WhatIsCLIPS.html#History

even the conventional World Wide Web!); nevertheless, CLIPS has been used in a number of studies to provide both DL reasoning and conventional inference capabilities with respect to OWL/RDF ontologies (Bassiliades & Vlahavas, 2004; Meditskos & Bassiliades, 2005; Meditskos & Bassiliades, 2006; Meditskos & Bassiliades, in press; Smart et al., 2007b).

The use of CLIPS as a DL reasoner, i.e. a system for computing RDF/OWL entailments, has been most extensively studied by Meditskos and colleagues. They have developed a system called O-DEVICE that imports OWL ontologies into CLIPS by transforming OWL constructs into an object-oriented model. O-DEVICE is an extension of a previous system, called R-DEVICE, which effectively maps RDF Schema constructs and data into CLIPS Object-Oriented Language (COOL) objects and then reasons over the RDF data using the CLIPS rule language (Bassiliades & Vlahavas, 2004).

CLIPS has also been used to provide a semantic reasoning capability with respect to Military Operations Other Than War (MOOTW) (Smart et al., 2007b). In this case, CLIPS was used to provide support for both DL reasoning and domain-specific decision-support. In particular, reasoning was directed towards the provision of decision support in three task areas:

- **Needs assessment:** an assessment of what needs to be done in terms of humanitarian relief actions in order to minimize further harm and alleviate human suffering.
- **Relief planning:** the actual planning of a relief effort in terms of the sourcing, delivery and dissemination of aid supplies.
- **Future vulnerability assessment:** an assessment, or prediction, of the long term implications of the disaster with respect to future humanitarian action.

The approach to ontology representation within this study was somewhat different to that employed by O-DEVICE. Most notably, a transformation process was employed in which the RDF triples of the OWL ontology were asserted into the CLIPS environment as simple fact assertions (see Figure 9). The advantage of this approach is that it avoids the complexity associated with the transformation of ontological statements into object-oriented constructs.

```
(triple
      (predicate "http://www.w3.org/1999/02/22-rdf-syntax-ns#type")
      (subject "http://sa.aktivespace.org/ontologies/aktivesa#IntelligenceMessageType")
      (object "http://www.w3.org/2002/07/owl#Class")
)

(triple
      (predicate "http://www.w3.org/2000/01/rdf-schema#subClassOf")
      (subject "http://sa.aktivespace.org/ontologies/aktivesa#IntelligenceMessageType")
      (object "http://sa.aktivespace.org/ontologies/aktivesa#Link16JSeriesMessageType")
)

(triple
      (predicate "http://www.w3.org/1999/02/22-rdf-syntax-ns#type")
      (subject "http://sa.aktivespace.org/ontologies/aktivesa#ChemicalWeapon")
      (object "http://www.w3.org/2002/07/owl#Class")
)

(triple
      (predicate "http://www.w3.org/2000/01/rdf-schema#subClassOf")
      (subject "http://sa.aktivespace.org/ontologies/aktivesa#ChemicalWeapon")
      (object "http://sa.aktivespace.org/ontologies/aktivesa#BiochemicalWeapon")
)

(triple
      (predicate "http://www.w3.org/1999/02/22-rdf-syntax-ns#type")
      (subject "http://sa.aktivespace.org/ontologies/aktivesa#AgentRelatedThing")
      (object "http://www.w3.org/2002/07/owl#Class")
)
```

Figure 9: RDF Triples as CLIPS Fact Assertions

2.3.3 Jess

The Jess[25] Java Expert System Shell (Friedman-Hill, 2000) is a popular rule system written in Java. The core Jess language is compatible with CLIPS; so many Jess scripts are also valid CLIPS scripts and vice-versa. In addition to the functionality supported by CLIPS, Jess incorporates JDBC technologies, which could potentially support a closer integration between the reasoning system and back-end knowledge repositories.

A number of studies have investigated the use of Jess as a rule engine for semantic reasoning (Golbreich, 2004; Golbreich & Imai, 2004; Grosof et al., 2002). Golbriech (2004) describes a mechanism for combining rule-based and DL semantic reasoners that utilizes the Jess and Racer reasoning systems respectively. Her strategy is to use Protégé-OWL as an editing environment for both the domain ontology and SWRL rule sets. The OWL ontology is then loaded into Racer, which automatically classifies the OWL classes and individuals using subsumption reasoning. Following this initial reasoning step, OWL instances are translated into Jess facts and SWRL rules are translated into Jess rules. The Jess Rule engine is then launched and any inferred facts are returned to Racer for the computation of additional semantic entailments. A loop is hereby established in which Racer and Jess are invoked in an iterative fashion until an inconsistency is detected (Racer) or no new fact is inferred (Racer and Jess) [26].

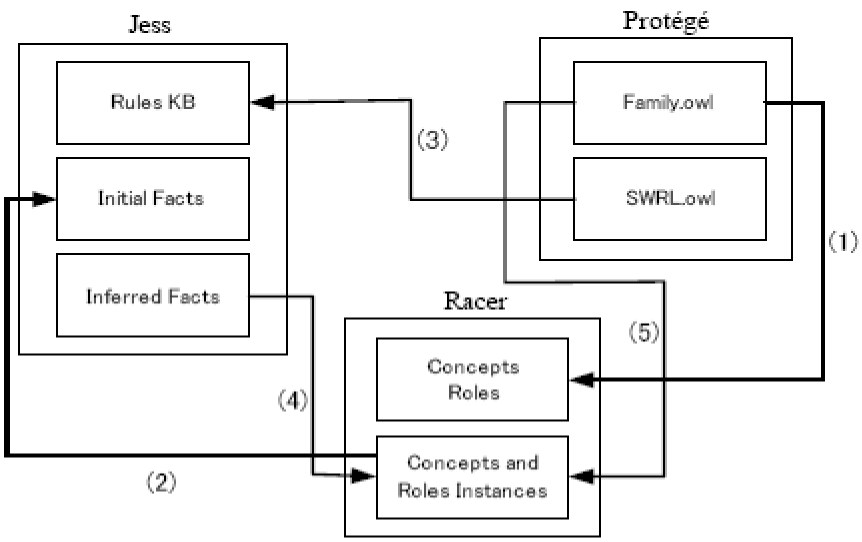

Figure 10: Using JESS and Racer for Combined Deductive and DL Reasoning

The combined deductive/DL reasoning capability described by Golbreich (Golbreich, 2004; Golbreich & Imai, 2004) is appealing, but the approach is not without its potential problems. Firstly, it is not clear how scalable the solution is. Benchmarking studies have suggested that the Jess system may not scale very well, and it seems to perform poorly relative to many other reasoning systems (at least when it is used to perform DL reasoning[27]) (Meditskos & Bassiliades, in press). In addition, any performance issues are likely to be exacerbated by the iterative processing of data between the

[25] http://herzberg.ca.sandia.gov/jess/

[26] This loop is established because any new facts asserted by Jess rules may support a series of further DL entailments.

[27] In some ways this is not a fair comparison because most DL reasoners are optimized, not surprisingly, for DL reasoning, whereas Jess provides a more general purpose reasoning capability.

Racer and Jess sub-components. The time taken to complete any reasoning task within this system is essentially the sum of:

1. the time taken to load the ontology into Racer
2. the total time taken by Racer to perform DL reasoning across successive processing cycles
3. the total time taken to transfer and translate OWL instances from Racer to Jess across successive processing cycles
4. the time taken to convert SWRL rules into Jess Rules
5. the total time taken by Jess to implement rule invocation and reasoning across successive processing cycles
6. the total time taken to transfer and translate Jess facts to Racer across successive processing cycles

This is a large overhead in terms of data transfer, translation and computation and one would expect performance to degrade significantly in situations where a large number of inferences or processing cycles are encountered. Nevertheless, Golbreich (2004) argues that an approach based on a combination of both DL and deductive reasoning is required for most semantic reasoning applications: inferences are essentially incomplete when either form of reasoning is used independently of the other.

Another system that relies on Jess to provide a semantic reasoning capability is ROWL[28]. ROWL enables users to create rules in RDF/XML syntax using the ROWL ontology. XSLT stylesheets are then used to transform the RDF/XML rules into forward-chaining rules using the Jess rule syntax. Additional stylesheets are used to transform the domain ontologies and instance data into unordered facts, each one representing an RDF triple (this is equivalent to the strategy described by Smart et al in the previous section).

2.3.4 Prolog

A number of rule engines are built on top of the Prolog programming language. One of these that has been studied in relation to semantic reasoning is SweetProlog (Laera et al., 2004). SweetProlog provides a system for translating OWL ontologies and rules (expressed in the proprietary rules language – OWLRuleML) into a set of Prolog programming constructs, specifically a set of facts and rules.

2.3.5 CWM

The Closed World Machine (CWM[29]) is an inference engine developed as part of the W3C Semantic Web Application Platform (SWAP[30]) initiative. It is essentially a forward chaining reasoner written in python that can be used for a number of general data processing tasks within the context of the Semantic Web, e.g. querying, checking, transforming and filtering information. Its core language is RDF, extended to include rules.

[28] http://www.cs.cmu.edu/~sadeh/MyCampusMirror/ROWL/ROWL.html
[29] http://www.w3.org/2000/10/swap/doc/cwm
[30] http://www.w3.org/2000/10/swap/

3 Military Application Areas

Given the need to maintain information and decision superiority in an era of network-enabled capabilities, it is perhaps not surprising that semantic technologies have been the considerable focus of attention in the military community. Specific applications of Semantic Web technologies can be found in areas such as information fusion (Eichmann, 1998; Kokar et al., 2004; Matheus, 2005; Scherl & Ulery, 2004; Sycara et al., 2003), coalition planning (Mott & Hendler, 2007), sensor selection (Preece et al., 2007) and modelling and simulation (Lacy & Gerber, 2004) to name but a few. In this section we explore a number of additional application areas for semantic technologies. These areas represent a subset of application areas where ontology-based reasoning capabilities could be particularly useful.

3.1 Semantic Integration & Interoperability

The ability to exploit semantically heterogeneous and physically disparate information sources is a key capability in terms of realizing the potential of large-scale networks to contribute to information superiority. As the global information environment becomes increasingly pervasive and spans ideologically, culturally and ethno-linguistically diverse communities, so the information exchange challenge for military agencies becomes ever harder. In future defence-related contexts, strategies for operationally-effective modes of information exchange and exploitation will need to target a variety of disparate information repositories and communication systems, including digital datalinks, military information repositories, and the totality of the information space available via internet-enabled, Network-Enabled Capability (NEC) and peer-to-peer computing environments. In situations such as these, the potential for semantic ambiguity is rife because the meaning of symbolic information often reflects the experiential, epistemic, cultural and task-specific biases of the information provider. Both the semantic referents and semantic significance of information is not invariant with respect to information exchange contexts, rather one sees a degree of semantic specificity - a community specific interpretation of meaning that may not necessarily transcend cultural, organizational and/or national boundaries. The point is that once we encounter distributed network environments that subtend a wide variety of information systems, sources and user communities, we face a critical challenge in terms our ability to integrate and share information in a semantically-sensible manner (one that respects the meaning assigned to information content by the originating agent or agency). Addressing the interoperability challenge is arguably one of the most important areas for semantic technology research in relation to future military capabilities (Smart & Shadbolt, 2007) and this is reflected in the attention given to semantic inter-operability issues by defence and government organizations, e.g. NATO recently established a Semantic Interoperability program specifically to investigate semantic interoperability issues in relation to heterogeneous C2 systems (Peter Houghton, personal communication).

At first blush it might seem as though reasoning technologies would be of minimal relevance to semantic integration and interoperability capabilities. A number of approaches to semantic integration have been studied by the research community (de Bruijn et al., 2006; Euzenat et al., 2004; Kalfoglou et al., 2005; Noy, 2004), and while all of these benefit from DL reasoning capabilities to some extent, none of them explicitly relies on conventional rule-based processing to effect information exchange and integration. A new approach that is being explored in the context of the joint US/UK International Technology Alliance (ITA[31]) initiative, however, does avail itself of rule processing capabilities. The approach is grounded in the notion that some forms of information exchange can be effected by SPARQL queries, in particular SPARQL queries that utilize a CONSTRUCT clause to translate knowledge statements from one ontology into (semantically-equivalent) statements in a second (related) ontology. As an example of this process, look at Figure 11. The SPARQL query illustrated in this figure converts instances of a 'Person' class (source:Person) in one ontology into instances of another 'Person' class (target:Person) in a separate ontology. The use of the SPARQL CONSTRUCT clause in this query exploits a mapping between the classes of the source ontologies to create a moderately complex information exchange solution. Note that the execution of the query does more than just instantiate instances of 'target:Person', it also transforms the '?carName' variable binding into instances of the 'target:Car' class and associates the new 'target:Person' and 'target:Car' instances with the 'target:drives' property.

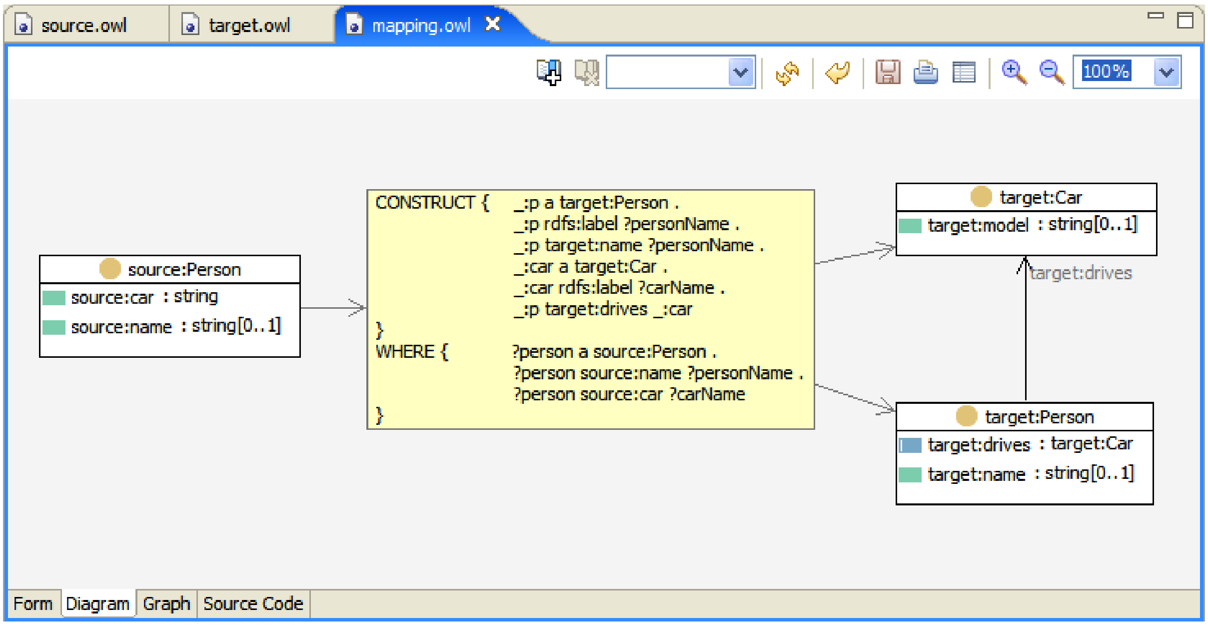

Figure 11: SPARQL Realization of Ontology Mapping Solution

As was mentioned in Section 2.1.1, SPARQL queries of this form can be viewed as deductive rules (Polleres, 2007) because the query is being used to derive new knowledge – in this case deriving new knowledge relative to a target ontology. Given that SPARQL CONSTRUCT queries are essentially rules, and that such queries can be used to implement semantic integration solutions, it may be worth examining more generic rule-based approaches in order to determine if advanced forms of semantic inter-operability could be supported by ontology-based reasoning. Our approach with respect to this issue, in the context of the ITA, is to examine whether rule-based solutions can be used as a generic mechanism to mediate information exchange between disparate agent communities, legacy applications and data repositories. One particularly interesting strand of research concerns the automatic instantiation of these reasoning solutions (in terms of automatic rule generation) based on the availability of semantically-rich annotation frameworks that are applied to the outcomes of ontology mapping/alignment solutions.

[31] http://www.usukita.org/

3.2 Decision Support

Conventional reasoning systems (those based on classical rule-based systems) have been used for many years to provide decision support capabilities in military contexts. Examples of such systems include mission planning (Smart, 2001; Smart, 2002), fighter pilot aiding (Shadbolt et al., 2000) and terrain analysis (Richbourg & Olson, 1996) systems. There are also cases where semantic reasoning is being used to provide military decision support. These include the Joint Explosive Ordnance Disposal Decision Support System (JEOD DSS) where OWL ontologies are being used to present relevant procedural information about EOD Tactics, Techniques and procedures (TTP) to warfighters (Lacy et al., 2005). The US Air Force (USAF) is also exploiting semantic capabilities to represent portions of the Foreign Clearance Guide (FCG) in support of automated mission planning capabilities (Lacy et al., 2005). According to Lacy et al (2005), one problem encountered with the FCG planning systems concerns the absence of an effective rule language to capture the kinds of expressions required by the planning system. He cites the following example of an FCG rule, which is difficult to represent using current (Semantic Web) rule languages:

> *"If a mission aircraft carries hazardous cargo and a country specifies that no mission carrying hazardous cargo can land, then each airbase associated with the country will not allow a mission carrying hazardous cargo to land." (Lacy et al., 2005)*

Clearly, an ability to develop sufficiently expressive rule languages and efficient semantic reasoners is a prerequisite for the successful application and acceptance of semantic technologies by military agencies.

3.3 Situation Awareness

Situation awareness is widely regarded as a critical enabler for many types of military operations; however, the modern military environment presents a number of challenges to situation awareness (Smart et al., 2007a). Firstly, it is worth noting that the introduction of advanced sensor systems and the pervasiveness of large-scale information networks are both a boon and a burden in terms of situation awareness. They are a boon inasmuch as such technologies provide us with an unprecedented opportunity to detect and communicate situation-relevant information, but they are a burden in the sense that end-users have to cope with an ever increasing quantity of information. A second issue concerns the impact of the new mission command philosophy, which places much more responsibility on junior commanders to possess both global and local situation awareness in order to make appropriate decisions (*The UK Joint High Level Operational Concept: An Analysis of the Components of the UK Defence Capability Framework*, 2005).

Figure 12: Contribution of Semantic Technologies to Enhanced Situation Awareness

There have been many studies investigating the contribution of semantic technologies to enhanced situation awareness (Matheus, 2005; Matheus et al., 2003a; Matheus et al., 2004; Matheus et al., 2005a; Matheus et al., 2005b; Matheus et al., 2003b; Smart et al., 2007a)[32] (see Figure 12). Perhaps most obviously, semantic technologies can improve the retrieval of relevant information from semantically heterogeneous and physically disparate information sources, thereby supporting information acquisition – effectively an increase in level 1 situation awareness (see Endsley, 2000; Endsley, 1995). They can also, however, exploit semantic reasoning capabilities to provide support for information triage, knowledge monitoring and information filtering capabilities (Smart et al., 2007b). For the most part, advanced knowledge monitoring and information filtering capabilities can be realized with a combination of DL reasoning and semantic queries; other types of reasoning, e.g. automated planning aids, may indirectly contribute to enhanced situation awareness by increasing attentional processing and working memory capacity.

3.4 Knowledge Discovery

The availability of a global repository of machine-interpretable information and data has enormous potential in terms of the realization of advanced knowledge discovery capabilities. The idea here is that data mining and rule induction techniques could be used to detect statistically significant patterns, associations and contingencies across large-scale, distributed information repositories. Consider the publication of multiple datasets pertaining to some of domain interest. Imagine for example, that several information providers publish information about the characteristics of various mushroom species. Once this information is represented in a form where a machine is able to correlate properties and interpret data in a semantically-coherent way, the potential for automatically discovering associations and dependencies within the datasets becomes a realistic possibility. It potentially enables automated processors to derive rules from datasets and represent these rules in a form that could be shared with, and exploited by, other agencies.

[32] see also http://eprints.ecs.soton.ac.uk/14351/2/B2_3.ppt

3.5 Semantically-Mediated Data/Information Fusion

The potential contribution of knowledge and semantic technologies to fusion-related problems has been recognized by a number of authors (Boury-Brisset, 2003; Matheus, 2005; Matheus et al., 2005b; Scherl & Ulery, 2004; Smart et al., 2005). While most analyses of semantically-mediated information fusion have concentrated on techniques for improved situation analysis, other analyses (e.g. Smart, 2005) have identified a plethora of opportunities at all levels of the JDL Data Fusion Model (see Llinas et al., 2004; Steinberg & Bowman, 2004).

One particular focus for analysis concerns the use of semantic technologies and reasoning capabilities in order to improve low-level entity recognition and feature extraction processes (Guo et al., 2007). While the ability to make semantic abstractions from low-level sensor data is largely regarded as beyond the current state-of-the-art (cf. Serre et al., 2007), it does appear that the exploitation of a variety of forms of contextual information, perhaps derived from a variety of disparate sources, could be used to improve object recognition. As an example of this process we are currently exploring the impact of contextual information on military vehicle recognition using microphone arrays and canonical acoustic profile datasets. The key research issue is whether information about factors such as the road surface type or meteorological conditions can be used to improve vehicle identification rates in realistic acoustic environments.

3.6 Machine-To-Machine Interaction

Semantic technologies can facilitate the interaction of automated agents and processors in order to yield problem-solving abilities that would be difficult to achieve by any one agent acting in isolation. While knowledge representation is one aspect of this ability, supporting the communication of semantically-coherent information, it is probably the case that reasoning systems will be required to assist Machine-to-Machine (M2M) interaction, if only for the purposes of establishing and coordinating agent collaborations. Stoutenburg et al (2005) additionally suggest that rule standardization is an important element of M2M capabilities:

> *"A standard rule framework built to operate over ontologies can enable machine to machine interfaces in a number of DoD environments. Dynamic C2 systems could employ M2M interactions for "asset allocation" of battlefield capabilities. For example, Unmanned Aerial Vehicles (UAV) could be dynamically tasked for reconnaissance and surveillance. Requests for backup and troop reallocation could also be automatically triggered during times of distress. Requests for clarifying information between components could be exchanged by applications while monitoring battlefield events, resulting in richer alerts and recommendations to the Warfighter." (Stoutenburg et al., 2005)*

3.7 Adaptive Information Flows

Managing information flows across the coalition battlespace is a complex problem. Specific problems arise when diverse sources provide C2 information in different formats, for different purposes with varying levels of security classification. Additional problems arise from the need to ensure that information is channelled to particular network nodes based on their task-specific information needs and concerns. Rules can be used to address some of these problems, ensuring the adaptive flow of information across a networked-environment, matching incoming information streams to particular consumers and actively transforming information content in ways that reflect the idiosyncratic processing capabilities of network nodes:

> *"…information sharing rules based on usage and capability of the receiving source can ensure that a particular node in the theatre receives only that information that it can process and display in a meaningful way…information sharing rules can be based on periodicity or events, so that updates will be sent to multiple partners on a dynamic basis." (Stoutenburg et al., 2005)*

4 Barriers to Technology Adoption

The core vision of the Semantic Web, as presented in Section 1, is clearly relevant to military capabilities, but what is the likelihood that this vision will actually materialize – are we just a few short steps from a network-enabled nirvana, or will we have to accept a more limited set of semantically-enabled capabilities? Clearly, the scope of this report prohibits a detailed analysis of the issues at stake here, but it is clear that some semantic technologies, at least in their current form, do possess some shortcomings. The extent to which these shortcomings undermine the possibility of early technology adoption by the defence community is the focus of the current section.

4.1 Open World Reasoning

Both RDF and OWL are specializations of predicate logic. In particular, they form a subset of predicate logic known as first order logic. First order logic is committed to the Open World Assumption (OWA), which asserts that a system's knowledge is incomplete. Under the conditions of the OWA, if a statement cannot be inferred from what is expressed in the system, then it still cannot be inferred to be false. If, for example, we have an ontology where the concept of a CivilianTerroristTarget is defined in terms of the logical complement of a MilitaryTerroristTarget concept (see Figure 13), then, under the conditions of the OWA, if a target (*X*) is not asserted to be an instance of a MilitaryTerroristTarget, then we cannot conclude that *X* is a CivilianTerroristTarget. The OWA is closely related to the monotonic nature of first-order logic: adding new information never falsifies a previous conclusion. If we learn at a later time that *X* is in fact a MilitaryTerroristTarget, then this conclusion does not invalidate the outcomes of any previous inferences (because in this case no inferences were made!). Reasoners that compute entailments from OWL ontologies, specifically DL reasoners such as Pellet and Racer, embrace the OWA and implement a form of reasoning called Open World Reasoning (OWR).

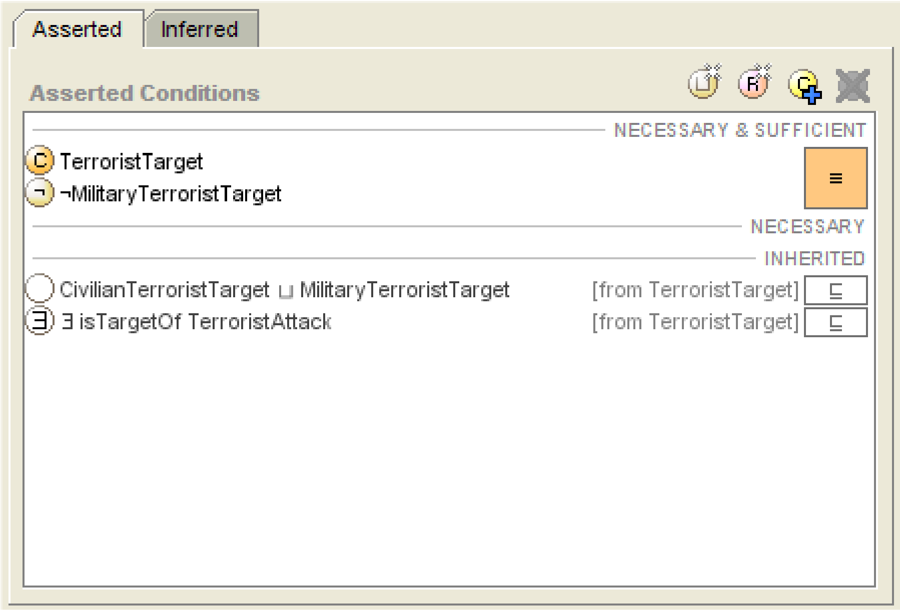

Figure 13: Definition of the CivilianTerroristTarget Concept[33]

In the context of the Semantic Web the OWA seems to be a reasonable assumption to make. The distributed nature of information content (in both space and time) means that we can seldom operate in a situation of complete information. In some situations, however, a strict commitment to OWR seems to be overly restrictive and it may depart significantly from much of the flavour of real-world reasoning and decision-making. Firstly, in many cases we are able to exploit background knowledge to make sensible guesses about the properties of things in the World. It is fair to assume, for example, that if *B* is Bird, then it can fly. This is despite the fact that there are clearly some instances of Birds that cannot fly. Since the OWA prohibits the assertion of statements that could be falsified in the light of additional information, a reasoner committed to the OWA would be unable to infer whether *B* can in fly or not. This form of 'decision paralysis' seems unsuitable in some situations[34], particularly those that require some form of immediate action (see also Section 4.2). Ideally, what is required here is some form of default reasoning whereby we can to state the value of missing information based on reasonable assumptions about the current situation and the statistical regularities that inhere in the target problem domain.

A corollary of the notion of default reasoning is the idea that, under certain conditions, if something has not been asserted then we can safely assume that it does not exist. Grimm and Motik (2005), for example, suggest the following:

> *"Consider a table of train departure times. If the table does not explicitly state that a train leaves at 12:47, then we usually conjecture that there is no such train. In other words, for train time-tables we typically use the closed world assumption (CWA), assuming that our knowledge about that part of the world is complete. Under CWA, we conclude that there is no train at 12:47 unless we can prove the contrary." (Grimm & Motik, 2005)*

---

[33] This example is taken from the Terrorism Ontology developed at the University of Southampton as part of the ITA initiative.

[34] Clearly there are some situations when the OWA seems perfectly acceptable and valid, e.g. situations when we need to err on the side of caution and cannot risk assuming something that has not already been asserted.

On occasion, a commitment to the OWA makes reasonable sense, especially in regard to the situations of incomplete information that one typically sees in distributed information environments. On the other hand, an overly stringent commitment to the OWA would seem to force us to renege on one of the most compelling features of real-world decision-making: the ability to make sensible assumptions about the nature of the World and coordinate response output accordingly. Inasmuch as OWR represents a significant departure from the form of reasoning that underpins decision-making in naturalistic settings, then it is unlikely that semantic reasoners can realistically expect to emulate the capabilities of a human reasoner. The danger, of course, is that these limitations will escalate to the point where they fall short of supporting the kind of capabilities expected by decision-makers in general, and military decision-makers in particular. What can we do to resolve the situation?

One potential solution strategy is to allow for a limited form of Closed World Reasoning (CWR) (reasoning the embraces CWR assumptions) in situations where the OWA is generally applicable. The Local Closed World Assumption (LCWA) (Doherty et al., 2000; Etzioni et al., 1997), for example, can be thought of as a useful compromise between the OWA and Closed World Assumption (CWA). Reasoning based on the LCWA allows us to augment OWA with the possibility of explicitly 'closing off' parts of the world in which the CWA applies. The application of the LCWA means that any query against a knowledge base (e.g. is *X* a MilitaryTerroristTarget?) will return the value 'true', 'false' or 'unknown'. With respect to our previous example in the domain of terrorist incidents, if *X* is a TerroristTarget, then we know that it must be either a MilitaryTerroristTarget or a CivilianTerroristTarget. This is because the semantics of the application domain dictate that these two classes are disjoint (one is the complement of the other) and they are complete with respect to the TerroristTarget class (see Figure 14). Now, if the OWA is in force, and we know that *X* is a TerroristTarget, but we cannot assert, for sure, whether *X* is a MilitaryTerroristTarget, then a reasoner will return the value 'unknown' in response to a query of the form '*is X a MilitaryTerroristTarget?*'. This contrasts with the value that will be returned if the CWA is in force. In this case, the reasoner will return 'false', because no information about *X* being a MilitaryTerroristTarget has been asserted in the knowledge base. Suppose we do know that there is only one MilitaryTerroristTarget in the world – MilitaryTerroristTarget(*Y*) and, therefore we know that *X* is not a MilitaryTerroristTarget (by implication it must be a CivilianTerroristTarget). How could we guarantee that our query correctly returns 'false' in response to the query '*is X a MilitaryTerroristTarget?*'? With the CWA in force this would already be the case as we have seen; however, with the OWA we would still have the value 'unknown' – not a particularly intelligent conclusion given what we know about the problem domain. The only way to correct this error would be to state the fact ¬MilitaryTerroristTarget(?) for the (potentially infinite) set of everything else in the world. This is clearly impractical; however, the LCWA overcomes the problem by allowing us to assert that the CWA applies in particular cases. McKenzie et al (2006) provide a practical demonstration of the LCWA in the context of a semantic reasoning system applied to the CS AKTive Space[35] (Shadbolt et al., 2004) domain. They demonstrate a LCWA solution based on the existence of two databases. The first database contains known facts describing the world, while the second database contains metadata indicating the categories of objects to which the CWA could be applied. McKenzie et al (2006) go on to demonstrate semantic reasoning capabilities that combine elements of both open and closed word reasoning. Such an approach has much to commend it in terms of

[35] http://cs.aktivespace.org/

implementing reasoning capabilities that possess more of the flexibility that characterizes human-level problem-solving in real-world environments.

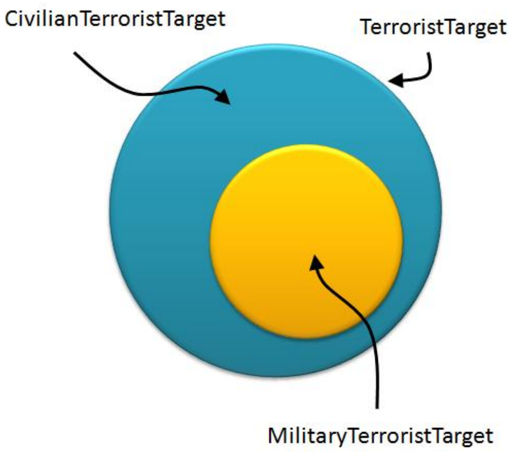

Figure 14: TerroristTarget Class

4.2 Uncertainty

In an uncertain world there are no absolutes! Uncertainty is pervasive when it comes to real-world decision-making. Seldom do we have complete information about a phenomenon of interest, and information gathering processes are seldom immune from inaccuracies or errors. The problem is particularly apparent in the case of the World Wide Web where inaccurate, incomplete and unreliable information is more or less commonplace. In addition to the uncertainty created by unreliable or incomplete information, there is also the uncertainty that derives from the inherent unpredictability of the world. We can seldom predict with absolute confidence the precise nature of future states the world, thus future plans and expectations are always somewhat provisional. This is partly what makes military planning such a difficult enterprise – it is genuinely difficult to anticipate enemy courses of action (in the case of warfighting operations) or the vagaries of the geological and meteorological environments in the case of humanitarian assistance missions. One is clearly reminded here of the old adage: “no plan survives contact with the enemy”.

The issue of uncertainty in the Semantic Web has been stressed multiple times (Kifer, 2005; Matheus, 2005; Stoilos et al., 2005), but one could be forgiven for embracing the rather pessimistic conclusion that we are still some way from a robust solution. Part of the problem, I think, lies in the commitment of semantic technologies to description logic formalisms. Such commitments make semantic systems inherently brittle and inflexible when it comes to real-world decision-making. When faced with a decision about whether a mushroom is poisonous or not, a system should be able to utilize available information to generate a decision – it should not simply conclude that it ‘does not know’ or wait until more information is forthcoming. In reality, situations of perfect information are seldom encountered, and when they are, the optimal time for selecting a response option has usually passed. Military commanders often (perhaps always) have to operate in situations of uncertainty and they have to make decisions within an operationally-effective timeframe, within a timeframe that maximally exploits temporary (and perhaps opportunistic) tactical advantages relative to enemy courses of action. If they are to do this, they need intelligent systems that can represent certainty and reason with uncertain information.

Subject	Predicate	Object
sem:MK-15_Type	rdf:type	sem:MineType
sem:MK-15_Type	sem:hasCountryOfOrigin	sem:USA
sem:MK-15_Type	sem:hasComponent	sem:SheetMetalCasingType
sem:MK-15_Type	sem:hasDetectability	sem:DifficultToDetect
sem:MK-15_Type	sem:UsesExplosive	sem:CompositionB
sem:MK-15_Type	sem:hasFragmentationRange	"500"^^xsd:int
sem:MK-15_Type	sem:hasColour	sem:Yellow
sem:MK-15_Type	sem:hasColour	sem:Olive
sem:MK-15_Type	sem:hasCountryOfUse	sem:Afghanistan
sem:MK-15_Type	sem:hasCountryOfUse	sem:Angola
sem:MK-15_Type	sem:hasCountryOfUse	sem:Cambodia
sem:MK-15_Type	sem:hasCountryOfUse	sem:Chile
sem:MK-15_Type	sem:hasCountryOfUse	sem:Cuba
sem:MK-15_Type	sem:hasCountryOfUse	sem:Cyprus

Figure 15: RDF Triples from the SEMIOTIKS Mine Action Ontology

Unfortunately, ontology representation languages such as OWL do not easily accommodate uncertainty information, and the logic systems on which they are based are not well-equipped to deal with uncertainty. One problem is that certainty information is notoriously difficult to represent within the representational framework of RDF. RDF, and its extensions, such as RDFS and OWL, is based on a scheme in which knowledge statements are expressed as triples. Figure 15 shows some example triples from one particular ontology that was developed in the context of the DIF DTC SEMIOTIKS[36] initiative. Each of these statements is assumed to be true, but what if we wanted to represent uncertainty information in this model, perhaps to reflect our degree of trust in the dataset provider? We cannot include an additional element in the triple to represent certainty information because this would violate one of the modelling assumptions of RDF. We could, however, use a mechanism whereby each subject is associated with additional predicates, each one of which represents the certainty associated with some other predicate/object pair (Jim Hendler, personal communication). For example, Figure 16 shows an approach to accommodating certainty information within RDF graphs that relies on the assertion of a datatype property (predicate) that is named after another property (associated with the same subject node) and distinguished from that property by the addition of a standard identifier (in this case the letter ‘P’). Let us call this the ‘P predicate strategy’. Unfortunately, this approach suffers from problems of application specificity – the interpretation of each ‘P predicate’ is local to the application and cannot be (automatically) understood by other, independent systems. Another problem relates to the ambiguity regarding triples with multiple predicates of the same type (see Figure 17). In this case, there is no way to relate individual ‘P predicates’ with their corresponding triples.

[36] http://www.edefence.org/semiotiks

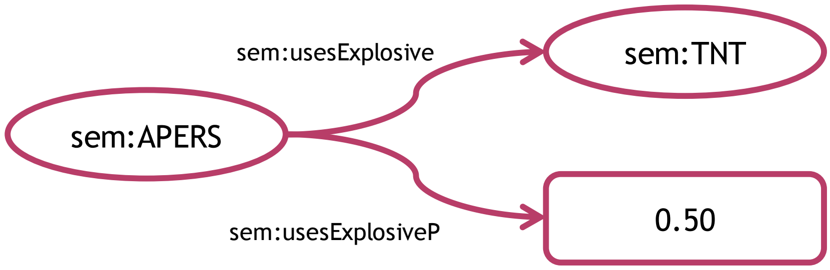

Figure 16: 'P Predicate' Approach to Representing Certainty Information

The development of certainty management solutions that enable existing representational formalisms and reasoning capabilities to deal with vague or imprecise knowledge is a key challenge that limits the applicability of semantic solutions to real-world decision-making contexts. If we cannot reliably represent certainty information for knowledge statements within the conventional representational framework of the Semantic Web, then it becomes difficult to assign probabilities or confidence limits to inference outcomes or the information items upon which such outcomes are based. In the absence of this key capability it is difficult to see how military agencies could wholeheartedly embrace semantic technologies as a generic solution for knowledge processing applications.

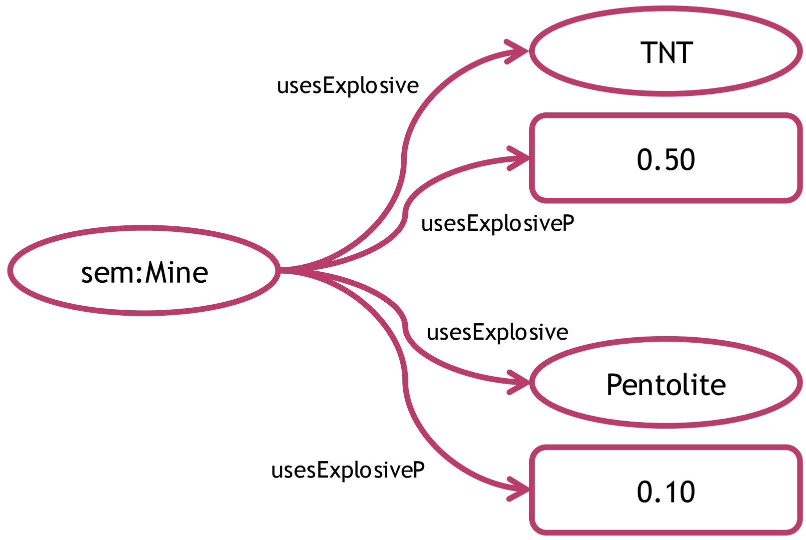

Figure 17: Breakdown of the 'P Predicate' Strategy

A number of strategies have been proposed to extend ontology representation and Semantic Web rules languages with an ability to represent vague or imprecise knowledge. These have largely assumed the form of probabilistic (Ding & Peng, 2004) and fuzzy logic extensions to OWL (Gao & Liu, 2005; Stoilos et al., 2006; Stoilos et al., 2005). Fuzzy logic formalisms are based on fuzzy set theory, which is a mathematical framework for covering vagueness (Klir & Yuan, 1994). Work in this area has given rise to Fuzzy-OWL (or f-OWL) (Stoilos et al., 2006; Stoilos et al., 2005), which can capture vague and imprecise knowledge. A Fuzzy Reasoning Engine[37], called FiRE, has also been developed, which lets Fuzzy-OWL capture and reason about uncertain knowledge (Stoilos et al., 2006).

In addition to fuzzy logic extensions to OWL, Pan et al (2005) have also proposed a fuzzy extension to SWRL. Their language, f-SWRL, avails itself of an ability to assign weights (in the range 0 to 1) to the

[37] http://www.image.ece.ntua.gr/~nsimou/

atomic formulas in both the antecedent (body) and consequent (head) of a SWRL rule. A rule in f-SWRL looks like something like the following:

$$A * \omega \leftarrow B_1 * \omega_1 \wedge \ldots \wedge B_n * \omega_n$$

where A, B_1, ..., B_n are either concepts (unary predicates) or properties (binary predicates) used in OWL DL axioms, and the weights ω_1, ..., ω_n and ω are real numbers in the unit interval.

4.3 Efficiency & Scalability

Efficiency and scalability are important characteristics of any reasoning capability, particularly those that are to be deployed in mission-critical military applications. The Semantic Web presents a number of challenges to scalable and efficient reasoning solutions, not least because the quantity of information made available by network infrastructures (and possibly the Semantic Web itself) may exceed the knowledge processing capacity of extant reasoning systems. Efficiency and scalability concerns are exacerbated by the relative richness of the semantic axioms used to encode domain-relevant knowledge. For example, the complexity of semantic reasoning for the different OWL variants is at least NP-hard (Hustadt et al., 2005), which indicates that, in general, semantic reasoning using OWL ontologies will not scale well. Furthermore, the semantic expressivity of ontologies may present problems from the perspective of efficient reasoning due to the large number of rule firings triggered by even relatively small changes to the underlying knowledge base (Smart et al., 2007b).

In their attempt to develop an efficient reasoning system for MOOTW operational contexts, Smart et al (2007b) provide a number of reasons to account for the poor performance of their reasoning solution. These include:

1. the time taken to retrieve, load and instantiate RDF triples from the knowledge repository into the reasoning environment, and
2. the performance overhead associated with rule execution.

Their proposed strategies for dealing with these problems include the following:

(i) optimization of the inference engine to support faster rule execution;
(ii) intelligent caching of temporary reasoning results;
(iii) progressive minimization of semantic expressivity (i.e. the removal of specific axioms) until a reasonable performance threshold has been attained;
(iv) modularization of ontology components to reduce inter-connectedness; and
(v) more precise control over the firing of specific rule subsets, i.e. only allowing certain rules to fire in a particular reasoning context.

On the basis of further research and reflection a number of other potential strategies can be proposed. They include:

(vi) delegation of reasoning tasks to optimized query engines;
(vii) optimization of the representational approach adopted for both ontologies and rules within specific reasoning environments;
(viii) use of approximate reasoning solutions;
(ix) optimal encoding of ontological data to support rule pattern matching; and

(x) incremental loading of ontological data and rule sets to reduce performance overheads.

In the remainder of this section we explore two potential strategies for delivering efficient reasoning solutions. These approaches represent a subset of those proposed in the above list.

4.3.1 Ontology Modularization

One way to improve the efficiency and scalability of reasoning systems is to minimize the number of knowledge statements and rule activations that need to be evaluated by the rule engine. If, for example, we could modularise an ontology to the extent where the effect of logical entailments (in terms of further rule activations) was limited to modular sub-components, then we could optimize rule processing by matching rules against a subset of the total number of knowledge statements contained in the global ontology. To make the example more concrete, suppose that a knowledge repository contains 3 ontologies (O1, O2, O3) and that each ontology is semantically-insular with respect to the other ontologies. This may happen in a situation where each ontology targets a separate domain and there are no overlaps between the domains at a conceptual level. Now suppose we have rules sets whose antecedent and consequent clauses are isolated to particular ontologies, say R1 is limited to O1, R2 to O2 and so on. In this case, we only need to consider O1 when a task (say T1) requires the evaluation of rules contained in R1. Furthermore, if changes are made to the ontology, then we only need to update the conclusions implied by the rules in R1 if the changes affect O1, otherwise we can safely assume that our previous conclusions are still valid. Using this strategy we could greatly limit the number of inference execution cycles that need to be undertaken by a reasoning system throughout its period of operational service. Of course, the extent to which we can realize this capability depends on our ability to partition ontologies into modular fragments and then link these components to specific rule sets that are themselves partitioned with respect to particular reasoning tasks (see Figure 18).

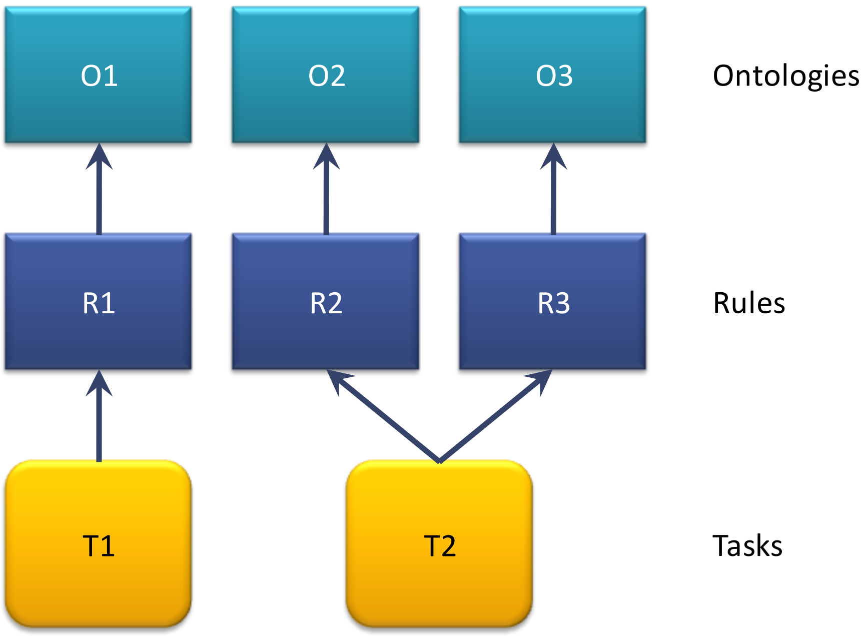

Figure 18: Ontology Modularization and Inference Optimization

While ontology partitioning capabilities have been investigated in previous studies (Stuckenschmidt & Klein, 2004), no studies have looked at the issue of ontology partitioning from the perspective of rule execution. Inasmuch as ontology modularization and ontology partitioning approaches constitute viable solution strategies for efficient reasoning systems, then it may be worthwhile considering modularization issues as a best practice principle with regard to ontology development, i.e. ontology developers should strive, as much as possible to create small modular ontologies with minimal degrees of interconnectedness. The extent to which this is possible in the Semantic Web, with its emphasis on distributed, yet heavily interconnected, knowledge resources is an issue that remains to be explored.

4.3.2 Incremental Loading of Ontological Data

In their attempt to develop efficient semantic reasoners within the CLIPS environment, Meditskos & Bassiliades (in press) describe an approach to rule and triple loading that, they argue, yields significant performance gains in terms of both rule execution and the computation of semantic entailments. Their approach is based on the incremental loading of triples (ILT) and rules (ILR) with intermediate bouts of rule activation and execution. Firstly with respect to ILR, they separate DL rules into 10 sets: transitive, symmetric, subproperty, inverse, equivalent, functional, inverse functional, universal quantifiers, existential and subsumption rules. Each rule set is loaded one at a time and all rule activations are processed before the next rule set is loaded. Meditskos & Bassiliades (in press) suggest this speeds up the pattern matching procedure to such an extent that it offsets the cost associated with iterative cycles of rule application:

> *"When the complete set of rules is loaded, the firing of one of them causes the pattern matching procedure to be executed over all rules in order to determine rule activations/ deactivations. By loading each time a portion of the rule set, the pattern matching procedure operates faster, even if the system spends extra time in order to apply the inference rules in a circular mode." Meditskos & Bassiliades (in press)*

A similar finding was observed with ILT. In this case the system incrementally loads sets of *q* triples, where *q* is a predefined value, and then applies the ILR methodology over the currently loaded dataset. The value of *q* has different effects on processing time depending on the specific ontology to be loaded. In particular, it has been argued that the value of *q* can be optimized by calculating a metric that represents the degree of complexity associated with ABox reasoning (Meditskos & Bassiliades, in press). Figure 19 illustrates the results obtained with different loading procedures against 6 different ontologies. The Direct Loading of Triples and Rules (DLTR) strategy corresponds to the direct loading of both rules and triples into the CLIPS environment in a single step, the ILR strategy corresponds to the incremental loading of rules strategy and the ILT+ILR corresponds to a combined strategy of incremental rule and triple loading. As is clear from this chart, ILT+ILR emerges as a more efficient solution as compared to either DLTR or ILR. If we focus specifically on the Lite-1 ontology, we can see that the ILT-ILR strategy has resulted in a 10-fold improvement in reasoning efficiency relative to the DLTR solutions (approx. 800 seconds vs. approx 8000 seconds respectively[38]).

[38] Actually, these figures should give us pause for thought because they highlight the scale of the efficiency challenges confronting us in terms of semantic reasoning capabilities. With the DLTR strategy and the Lite-1 ontology, it takes approximately 8000 seconds (or 2.22 hours) to complete rule processing.

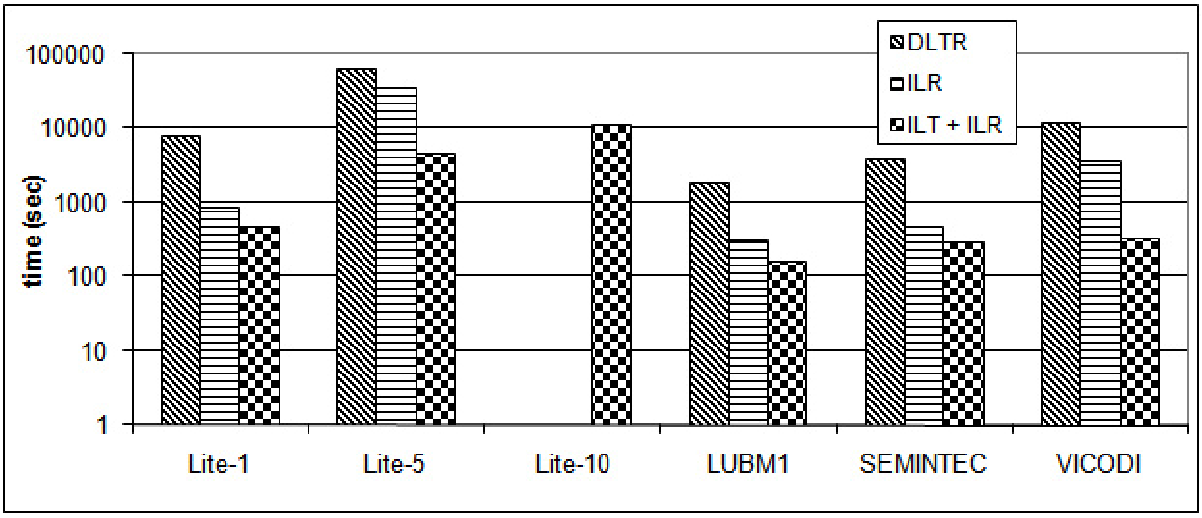

Figure 19: Effect of Different Loading Solutions on Reasoning Efficiency

4.4 Explanation Facilities

The ability to provide explanatory information is the hallmark of cognitively-transparent, symbolic knowledge-based systems and it allows the user (both expert and non-expert alike) to evaluate the logic of reasoning processes underlying a particular decision outcome. The value of explanatory capabilities with respect to semantic reasoning, even DL reasoning, has been emphasized a number of times in the literature, most often in terms of supporting the end-user with respect to the interpretation, comprehension and evaluation of logical inferences:

> *"…one would like to see an explanation for unintended specialisation links or inconsistencies and be supported in debugging such flaws; one would like to extract, from an ontology, the 'sub-ontology' that 'covers' a given term or concept; one would like to ask the editor to propose a new concept description as (the most specific) generalisation of a given set of instances; one would like to find a concept description that follows a certain 'pattern' of a concept; or one would like to see a user-friendly approximation of a concept description, for instance in a frame-based notation." (Sattler, 2007)*

There is no doubt that explanation facilities would of tremendous value in terms of understanding the logical (and, on occasion, the seemingly illogical!) entailments of DL reasoning. In fact, one proposal argues for the development of explanatory systems that can be utilized by other automated reasoners to diagnose and remedy faults (e.g. logical inconsistencies) with complex domain ontologies (Baclawski et al., 2004).

Explanatory capabilities are an essential element of reasoning systems in the military domain. The nature of military operations means that human decision-makers must be able to evaluate the reasons why a system reached a particular conclusion and the inferential strategy it employed to do so. The explanations provided by a system can influence the extent to which decision outcomes are sanctioned or countermanded by a military commander and it is therefore imperative that a reasoning system is able to communicate the reasons and rationale underpinning its inference steps in a way that is understandable (and acceptable) to human end-users. Clearly, this argues for the development of explanatory capabilities that are grounded in the medium of natural language, but there are a number of other factors that influence the comprehensibility (and ultimately the acceptability) of a reasoning process. One factor is the extent to which explanations are couched in the vocabulary of the target knowledge domain – with its idiosyncratic terms and phrases, the military domain makes vocabulary specialization more or less mandatory. Another factor, concerns the expertise level of the human agent – explanations that are suited for an experienced commander are not necessarily suited for someone at a lower expertise level, perhaps an initiate or trainee commander. Similar arguments can be espoused in the case of task focus, perspective and position in command hierarchies – a high-level commander does not necessarily require the same kind of justifications as a subordinate commander or commander at a different echelon of command, nor is the same level of detail with respect to explanations/justifications required. Related to these concerns are issues about the time available for evaluating decision outcomes with respect to machine-generated explanations. In most contexts, the generation of reams of detailed information about the inner workings of a reasoning system are unlikely to be of much use. In the case of the cognitive cockpit program (Shadbolt et al., 2000), for example, decisions about the use of a defensive aid suite in aerial combat situations had to be made on the order of a few seconds. In this case, there was simply no time to evaluate decisions in terms of system-generated explanations[39].

It is important to bear in mind that comprehensibility and acceptability are not the same thing. Just because a reasoning solution is presented in terms that are comprehensible to a military commander does not mean that it will make sense in relation to the assumptions, explanations and rationalizations that are typically employed by him/her to countenance particular decision alternatives. Unless a commander can interpret the 'logic' of a reasoning process with respect to his/her own systems of justification and rationalization, then the acceptability of an inference outcome (and the reasoning system in general) is likely to be undermined , at the very least it may be treated with some suspicion. This is a potential problem when it comes to semantic reasoning (and logical reasoning, in general) because the types of explanations that may be supported by logical entailments are not necessarily the kinds of explanations that human problem-solvers expect to see. First of all, human problem solvers may have developed reasoning strategies and explanatory styles that are largely domain-specific and acquired through years of experience and training. Furthermore, human decision-making is subject to a variety or 'errors' and biases, including (but certainly not limited to[40]) confirmation bias (Wason, 1960), fundamental attribution error (see Augoustinos et al., 2006) and anchoring (see Augoustinos et al., 2006). Sometimes these biases may conspire to undermine the acceptability or 'sensibility' of machine-generated reasoning outcomes. Some commentators have even suggested that irrationality is the hallmark of human decision-making (Sutherland, 1994), something which if true would certainly conflict with the obsession of the description logic community with logically valid entailments. In any case, it certainly appears that human beings experience difficulty in dealing with conventional logic problems (see Eysenck & Keane, 1995), and on occasion their reasoning is neither logical nor rational.

In summary then, we can see that semantic reasoning systems will, in all likelihood, need to avail themselves of explanatory facilities, and such facilities are likely to be mandatory in the military domain. The current analysis (which is grounded in our previous experiences of developing military

[39] Although, such information was used for system evaluation, debugging and training purposes.
[40] see http://en.wikipedia.org/wiki/List_of_cognitive_biases for a comprehensive list

decision support systems) suggests that explanation facilities need to consider a number of requirements, most of which go beyond the remit of typical expert system development initiatives – they certainly exceed anything that has been attempted in the context of the Semantic Web. Military semantic reasoning systems should, in general, satisfy the following requirements:

1. They should use a communication medium that is easily understandable to the human end-user, preferably the medium of natural language.
2. They should attempt to use domain-specific terms and phrases as much as possible.
3. They should adjust the type of explanation, and the level of detail given, to reflect the perspectives, needs and requirements of the human decision-maker.
4. They should only present as much information as is required to enable a human decision-maker to evaluate the integrity of the inference outcome.
5. They should generate detailed traces of inference execution in associated with explanatory information for the purposes of offline analysis, system validation and training.
6. They should consider the types of explanations given for actions and events in the domain. In particular, they should aim to follow the same 'logic' as used by human problem solvers in reaching a decision within the target domain. This will generally focus on domain-specific problem-solving strategies, but it may also draw on aspects of cultural and cognitive psychology as well as social cognition.

This is not intended as an exhaustive list, but it does at least highlight some of the challenges confronting us in terms of the need to develop better explanation facilities for semantic reasoning systems. Above all, the analysis suggests that greater attention needs to be paid to the domain-specific aspects of reasoning and decision-making[41]. It is imperative that we focus our attention on the idiosyncratic features of reasoning within particular task contexts and that we consider the cultural, psychological and social cognitive aspects of human decision making within these contexts.[42]

[41] Although it does not eschew the notion of developing more general purpose systems which could be valuable in certain contexts, e.g. ontology development, system evaluation, knowledge validation, performance optimization and 'debugging' of erroneous inference processes.

[42] A number of knowledge engineering techniques have been developed to support the analysis of task- and domain-specific knowledge (Schreiber et al., 2000; Shadbolt & Burton, 1990; Shadbolt et al., 1999). It is clearly worth investigating the application of these techniques to address the task and knowledge analysis requirements being proposed here. This includes the accounts that people provide for the events, decisions and actions that occur within the domain as well as the 'logic' that underpins domain-specific reasoning

5 Timeline for Technology Adoption

Section 4 provided an overview of technology barriers that need to be overcome in order to promote the uptake of semantic technologies by defence organizations. The current section builds on this analysis by proposing a timeline for technology adoption. The starting point for our analysis is a consideration of the military application areas presented in Section 3. We propose a timeframe for the uptake of semantic technologies in each of these capability areas by assessing the current state-of-the-art, the rate of research progress and the investment of government and defence organizations in semantic technologies.

5.1 The Evolutionary Timeframe of the Semantic Web

The Semantic Web initiative really began to gain momentum as a research program with the publication of Tim Berners-Lee et al's publication about the Semantic Web in 2001 (Berners-Lee et al., 2001). Since then, a number of standards have been established and there is a considerable degree of maturity with respect to semantic technologies, particularly in the areas of ontology development, semantic querying and knowledge storage (Shadbolt et al., 2006). This maturity is reflected in the growing number of commercial vendors providing support for semantically-enabled capabilities. Examples include Adobe's RDF based Extensible Metadata Platform (XMP)[43], Google's planned search engine enhancements through the acquisition of Applied Semantics[44], IBM's emerging Semantic Web platform[45] and Oracle's Semantic Technologies Centre[46]. Despite the progress, however, it is unclear at what point the Semantic Web vision, as proposed by Berners-Lee et al (2001) will be realized. One proposal suggests that the timeframe for wide-scale exploitation of semantic technologies on the Semantic Web will probably occur around the 2010/2011 timeframe (Berners-Lee & Hendler, 2001) (see Figure 20). This proposal, however, does not discriminate between different semantic technologies and capabilities that may have very different evolutionary timelines. According to Pulvermacher et al (2005) there are likely to be a number of short-term routes to technology adoption, particularly in the military domain. These include Semantic Web services, semantically-aware searches, semantically-guided information retrieval, and the provision of common ontologies to promote coalition inter-operability. In our view, this is a credible list of capabilities that could be adopted and deployed by military agencies within the 2010/2011 timeframe. The exploitation of domain ontologies for the purposes of information interchange is likely to be the most common use of semantic technologies within this timeframe, a sentiment that is shared by Semy et al (2004). Writing in 2004, they suggest that within the 2009 timeframe we will witness an increasing move towards the development of mid-level ontologies in domains such as command and control, operations, intelligence and logistics; an increasing move towards ontology modularization; and support for the exploitation of ontologies via automated discovery, registration and mapping process.

[43] http://www.adobe.com/products/xmp/index.html
[44] http://www.google.com/press/pressrel/applied.html
[45] http://www.alphaworks.ibm.com/topics/semantics
[46] http://www.oracle.com/technology/tech/semantic_technologies/index.html

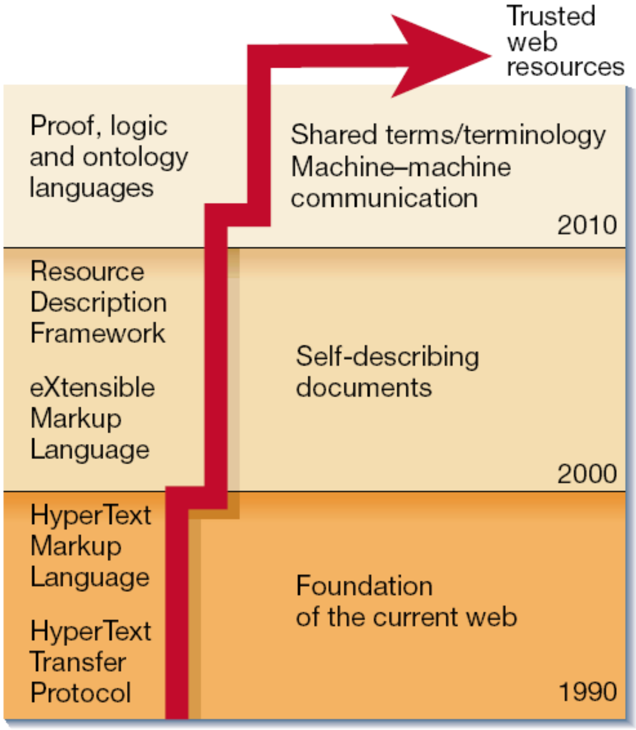

Figure 20: Timeframe for the Evolution of the Semantic Web[47]

5.2 The Emergence of Semantically-Mediated Reasoning Capabilities

Unlike the capabilities based on relatively mature semantic technologies (ontologies, semantic queries, knowledge repositories, and the like), the timeframe for capabilities based on semantic reasoning technologies is, we suggest, likely to be significantly longer than the 2010/2011 timeframe. This reflects the need for further research and development in a number of areas, including: efficient and scalable reasoning engines, rule language standards, inter-operability specifications, techniques for rule elicitation and acquisition, certainty representation and management solutions, approximate reasoning capabilities, human-oriented explanation facilities, and an ability to combine OWR with CWR techniques. Of course, not all semantically-enabled capabilities rely on the resolution of *all* these problems in order to demonstrate tangible benefits. Based on the analysis of military application areas presented in Section 3 we propose the progressive realization of reasoning-related capabilities according to the timeframe presented in Figure 22. While this proposal is, of course, speculative (and therefore provisional), it reflects, we feel, a plausible timeframe for technology adoption and exploitation based on a number of assumptions. As a means of explicating these assumptions, Figure 21 presents an evaluation matrix that rates each of the military application areas (identified in Section 3) against a number of evaluative criteria. These criteria include:

1. **Technological Maturity:** the extent to which the application area depends on mature semantic technologies, such as ontologies and semantic queries, as compared with relatively immature technologies, such as efficient reasoning engines. A high score on this dimension reflects a high level of dependency on mature technologies.

[47] source: Berners-Lee & Hendler (2001)

2. **Research Investment:** the extent to which the application area is the focus of current research programmes and technology development initiatives. A high score here reflects a high level of investment in research by government, commercial and research organizations.
3. **Military Impact:** an estimate of the potential impact of the application area on military capabilities, e.g. enhanced situation awareness may deliver more of an impact than domain-specific decision support systems. A higher score reflects a larger relative impact.
4. **Technical Feasibility:** the difficulty associated with providing technological solutions that support the desired capability. High scores reflect more feasible or easier solutions.

Capability Area	Technological Maturity	Research Investment	Military Impact	Technical Feasibility	Total
Semantic Integration & Interoperability	4	4	5	4	17
Decision Support	1	2	2	2	7
Situation Awareness	3	5	5	4	17
Knowledge Discovery	2	1	4	3	10
Data Fusion	4	5	4	4	17
M2M Interaction	3	3	3	3	12
Adaptive Information Flows	3	3	4	3	13

Figure 21: Evaluation of Semantically-Enabled Capabilities

In completing the analysis each capability area was evaluated with respect to the criteria using a five-point scale (see Figure 21). Given that high scores in this matrix are likely to reflect the possibility of early technology adoption, we can see that technologies supporting data fusion, situation awareness and semantic integration capabilities are likely to be adopted earlier than those associated with knowledge discovery and domain-specific decision support. Based on the current maturity of technologies supporting 'early adoption' applications, we would expect to see the widespread use of light-weight reasoning system solutions (underpinning these applications) within the 2015 timeframe (see Figure 22). In contrast, we estimate that robust decision support systems grounded on a Semantic Web infrastructure are unlikely to be in widespread use much before 2025.

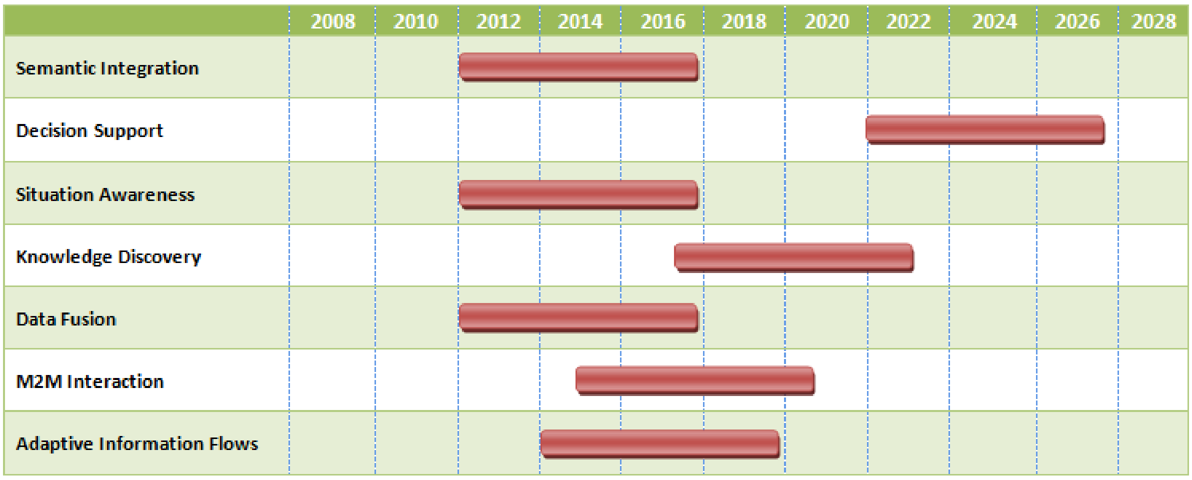

Figure 22: Timeline for Technology Adoption

6 Conclusion

This report has sought to review research relating to the development of semantic reasoning capabilities within the military domain. In general, it would seem that semantic technologies are relevant to future military capabilities, largely because they provide a platform for advanced information exploitation in a distributed network environment. This speaks directly to the military's need for information and decision superiority, a capability that will be grounded in the ability to identify, retrieve and integrate relevant information in an accurate and timely manner. Semantic technologies such as ontologies, semantic queries and knowledge repositories are already approaching a level of maturity where they could be deployed alongside existing military information systems. Large scale research programs such as the DIF DTC and ITA are already contributing to the early deployment of semantic technologies to front-line users with the roll-out of a proposed semantic solution to Afghanistan in July 2008. Inasmuch as the technological environment of the public, conventional Web matches the features of military network environments[48], then semantic capabilities demonstrated in the context of the conventional Web would be expected to generalize well to the military domain.

In contrast to ontologies and semantic queries, reasoning systems and rule languages for the Semantic Web are at an early stage of development. This is reflected in the lack of any consensus about the kind of rule language that should be adopted for the semantic web (although the standardization effort established by the W3C is likely to resolve this issue within the next couple of years) and the absence of any general purpose, high performance rule engines (academic or commercial) that can support advanced knowledge processing. Some progress has undoubtedly been made with respect to the characterization of performance deficits in existing systems, and there have been some promising developments with respect to the explication of performance optimization strategies in recent months (see Section 4.3). Nevertheless, it is probably safe to conclude that the existing state-of-the-art with respect to semantic reasoning capabilities is inadequate in many respects, particularly with respect to synthetic knowledge-intensive task types (see Schreiber et al., 2000) which may involve millions of knowledge statements and a similar number of rule activations. Of course, such conclusions only apply to reasoning processes of considerable complexity, and it is clear that not all tasks may incur the same performance overhead as those seen in situations like humanitarian relief planning (see Smart et al., 2007b). One area where current capabilities may be sufficient is in the area of semantic integration and interoperability. Semantic integration has been identified as a key capability for the defence community, promising to resolve some of the issues surrounding information exchange and the shortcomings of conventional data exchange solutions, such as Information Exchange Data Models (IEDMs) (Smart & Shadbolt, 2007). There is a widespread recognition of the value of semantic integration capabilities within the defence community - the recent creation of a Semantic Integration working group within NATO is testimony to this conclusion. Experimental analyses

[48] This may be a somewhat optimistic assumption. One difference between the conventional internet and emerging military networks concerns the increasing emphasis on wireless, ad hoc, and mobile capabilities. This shift in emphasis is not, of course, restricted to military contexts - mobile and wireless devices are pretty much ubiquitous nowadays - the problem is that it is not entirely clear whether a set of technologies that were developed in the context of the conventional internet are still suited to network architectures that violate some of the assumptions on which the internet is based.

undertaken in the context of the ITA initiative, suggest that semantic integration solutions grounded in rule-based processing may not suffer from the same kind of performance overheads as their decision-support counterparts (largely because of the relatively simplicity of the inference steps being performed). As such, semantic integration and inter-operability capabilities may be a useful target area for technology adoption by the military community with potential exploitation opportunities clearly defined at the national (UK Army) and international (e.g. NATO) levels.

Another barrier to technology adoption concerns the development of adequate explanation facilities (see Section 4.4). There are no semantic reasoning systems at present (that we are aware of) that would satisfy the criteria for explanatory support as outlined in this report. Nevertheless, the technological barriers to be surmounted in developing an explanatory capability are relatively minor, notwithstanding the requirement to consider the psychosocial and socio-technical context in which such a capability is to be deployed.

Unlike the issues surrounding the development of efficient and scalable reasoning systems, and the development of integrated explanation capabilities, the issues surrounding open world reasoning (see Section 4.1) and uncertainty management (see Section 4.2) are somewhat more problematic. The concern here is that the assumptions and commitments made by the Semantic Web community at the outset of the Semantic Web initiative may not be particularly well-suited to the realization of an effective, human-oriented assistive intelligence capability. Of particular concern is that the representational commitments made by the Semantic Web, largely by virtue of its origins in description logic, now negate the implementation of flexible modes of reasoning that resemble human-level reasoning in naturalistic settings, especially with respect to the capacity for graceful degradation and default reasoning as well as the ability to deal with incomplete and uncertain information, fuzzy concept categories and logical inconsistencies. There is, to some extent, a general feeling of unease that semantic reasoning systems fall foul of the same set of criticisms (e.g. excessive rigidity and lack of insight) levied against their conventional AI forbears a generation ago (Dreyfus, 1981). It is much more difficult here, I think, to see a way forward, in part because the problems are not so much technological as a combination of philosophical (what are the essential ingredients of an intelligent system?) and political (what kind of Semantic Web do we want, and who is best placed to deliver it?). There are a number of options in terms of future work (not all of which entail the preservation of the intellectual status quo!); however, the strategy that is most likely to succeed in the short term is to realistically assess the kinds of situations in which current reasoning capabilities are best placed to deliver added value. Our early experiences and experimental results suggest that the most acceptable (and perhaps the most useful) solutions from the user's perspective are those in which reasoning capabilities do not attempt to emulate human-level reasoning. Instead, perhaps the best opportunities for technology adoption (at least *early* adoption) lie in respect of 'light-weight' reasoning processes that complement and build on existing Semantic Web capabilities, e.g., the potential for enhanced search and retrieval capabilities, the ability for advanced modes of information exchange and integration, and the aggregation of task-relevant information for the purposes of enhanced situation awareness.

This report has, for the most part, emphasized the challenges posed by the military environment – the factors that may hinder or impede the whole-hearted adoption of semantically-enabled capabilities within the defence community. It is worth concluding, I think, by emphasizing the opportunities presented by the military domain in terms of the continued development and growth of the Semantic Web. Part of the problem with the Semantic Web is that its full potential is only fully realized when a certain critical mass has been obtained – when there are a sufficient number of technology adopters contributing to the progressive growth of a large-scale, semantically-enriched, global knowledge repository. In this sense, the military domain is an attractive proving ground for the Semantic Web. It provides an opportunity for practical demonstrations of semantic capabilities within an organizational context where technological innovation and change is both practical and desirable. The military community strives for capabilities that supersede those within easy reach of its adversaries. The Web as we know it is easily accessible to more or less anyone and it provides the means for adversaries to challenge the potential technological supremacy of the military, at least in respect of future network-enabled capabilities. In this sense, the adoption of Semantic Web technologies may be a necessary step in ensuring the future information and decision superiority of our military forces. It is obviously important to consider the benefits to the military of adopting a technology like the Semantic Web. Of equal importance, however, is the need to consider the potential benefit to our adversaries if we fail to fully realize the transformative potential of the Semantic Web.

7 References

1. Allemang, D. (2006). Rule-based intelligence in the Semantic Web-or-"I'll settle for a web that's just not so dumb!" *2nd International Conference on Rule Markup Languages for the Semantic Web*. Athens, Georgia, USA.
2. Antoniou, G., & van Harmelen, F. (2004) Web Ontology Language: OWL. In S. Steffen & R. Studer (Eds.), *Handbook on Ontologies*. Berlin: Springer Verlag.
3. Antoniou, G., & Wagner, G. (2003). Rules and Defeasible Reasoning on the Semantic Web, *2nd International Workshop on Rules and Rule Markup Languages for the Semantic Web*. Sanibel Island, FL, USA.
4. Augoustinos, M., Walker, I., & Donaghue, N. (2006) *Social Cognition: An Integrated Introduction* (2nd ed.). London, UK: Sage Publications Inc.
5. Baader, F., & Sattler, U. (2001) An overview of tableau algorithms for description logics. *Studia Logica, 69*(1), 5-40.
6. Baclawski, K., Matheus, C., Kokar, M., Letkowski, J., & Kogut, P. (2004) *Towards a symptom ontology for semantic web applications.* Paper presented at the 3rd International Semantic Web Conference, Hiroshima, Japan.
7. Bassiliades, N., & Vlahavas, I. (2004). R-DEVICE: A Deductive RDF Rule Language, *3rd International Workshop on Rules and Rule Markup Languages for the Semantic Web (RuleML2004)*. Hiroshima, Japan.
8. Berners-Lee, T., & Hendler, J. (2001) Publishing on the semantic web. *Nature, 410*(6832), 1023-1024.
9. Berners-Lee, T., Hendler, J., & Lassila, O. (2001) The Semantic Web. *Scientific American, 284*(4), 34-43.
10. Berstel, B., Bonnard, P., Bry, F., Eckert, M., & Patranjan, P. L. (2007). Reactive Rules on the Web, *Reasoning Web*. Dresden, Germany.
11. Boley, H., Kifer, M., Patranjan, P. L., & Polleres, A. (2007). Rule Interchange on the Web, *Reasoning Web*. Desden, Germany.
12. Boley, H., Tabet, S., & Wagner, G. (2001). Design Rationale of RuleML: A Markup Language for Semantic Web Rules, *International Semantic Web Working Symposium (SWWS)*. California, USA.
13. Boury-Brisset, A.-C. (2003) *Ontology-based approach for information fusion.* Paper presented at the 6th International Conference on Information Fusion, Cairns, Queensland, Australia.
14. Bry, F., & Marchiori, M. (2005). Ten Theses on Logic Languages for the Semantic Web, *3rd International Workshop on Principles and Practice of Semantic Web Reasoning (PPSWR 2005)*. Dagstuhl Castle, Germany.
15. Clark, A. (1989) *Microcognition: Philosophy, Cognitive Science, and Parallel Distributed Processing*. Cambridge, Massachussets, USA: MIT Press.
16. de Bruijn, J., Ehrig, M., Feier, C., Martin-Recuerda, F., Scharffe, F., & Weiten, M. (2006) Ontology mediation, merging and aligning. In J. Davies, R. Studer & P. Warren (Eds.), *Semantic Web Technologies : Trends and Research in Ontology-Based Systems*. Chichester, UK: John Wiley & Sons Ltd.
17. Ding, Z., & Peng, Y. (2004). A probabilistic extension to ontology language OWL, *37th Hawaii International Conference on System Sciences (HICSS-37)* (pp. 111-120). Big Island, Hawaii.
18. Doherty, P., Lukaszewicz, W., & Szalas, A. (2000). Efficient Reasoning Using the Local Closed-World Assumption, *9th International Conference on Artificial Intelligence: Methodology, Systems, and Applications (AIMSA 2000)*. Varna, Bulgaria.
19. Dreyfus, H. (1981) From micro-worlds to knowledge representation: AI at an impasse. In J. Haugeland (Ed.), *Mind Design* (pp. 161-205). Massachusetts, USA: MIT Press.

20. Eichmann, D. (1998) *Ontology-Based Information Fusion.* Paper presented at the International Conference on Intelligent User Interfaces: Workshop on Real-Time Intelligent User Interfaces for Decision Support and Information Visualization, San Francisco, CA.
21. Endsley, M., R. (2000) Theoretical underpinnings of situation awareness: A critical review. In M. R. Endsley & D. J. Garland (Eds.), *Situation Awareness Analysis and Measurement*. Mahwah, NJ: Lawrence Erlbaum Associates.
22. Endsley, M. R. (1995) Toward a theory of situation awareness in dynamic systems. *Human factors, 37*(1), 32-64.
23. Etzioni, O., Golden, K., & Weld, D. S. (1997) Sound and efficient closed-world reasoning for planning. *Artificial Intelligence, 89*(1-2), 113-148.
24. Euzenat, J., Le Bach, T., Barrasa, J., Bouquet, P., De Bo, J., Dieng, R., Ehrig, M., Hauswirth, M., Jarrar, M., Lara, R., Maynard, D., Napoli, A., Stamou, G., Stuckenschmidt, H., Shvaiko, P., Tessaris, S., Van Acker, S., & Zaihrayeu, I. (2004) *State of the art on ontology alignment.* Karlsruhe, Germany: University of Karlsruhe. (D2.2.3)
25. Eysenck, M. W., & Keane, M. T. (1995) *Cognitive Psychology: A Student's Handbook* (3rd ed.). Hove, UK: Lawrence Erlbaum Associates.
26. Friedman-Hill, E. J. (2000) *Jess, The Java Expert System Shell.* Livermore, CA: Sandia National Laboratories.
27. Gao, M., & Liu, C. (2005). Extending OWL by Fuzzy Description Logic, *17th IEEE International Conference on Tools with Artificial Intelligence* (pp. 562-567). Hong Kong: IEEE Computer Society.
28. Giarratano, J. C., & Riley, G. (1994) *Expert Systems: Principles and Programming*. Boston, MA, USA: PWS Publishing Co.
29. Golbreich, C. (2004). Combining Rule and Ontology Reasoners for the Semantic Web, *3rd International Workshop on Rules and Rule Markup Languages for the Semantic Web (RuleML)*. Hiroshima, Japan.
30. Golbreich, C., & Imai, A. (2004). Combining SWRL rules and OWL ontologies with Protégé OWL Plugin, Jess, and Racer, *7th International Protégé Conference*. Bethesda, Maryland, USA.
31. Grimm, S., & Motik, B. (2005). Closed World Reasoning in the Semantic Web through Epistemic Operators, *1st International Workshop on OWL Experiences and Directions*. Galway, Ireland.
32. Grosof, B. N., Gandhe, M. D., & Finin, T. W. (2002). SweetJess: Translating DamlRuleML to Jess, *International Workshop on Rule Markup Languages for Business Rules on the Semantic Web (RuleML 2002)*. Sardinia, Italy.
33. Guo, B., Wang, Y., Smart, P. R., Nixon, M., Shadbolt, N. R., Preece, A., Damarla, T., & Brown, M. (2007) *Approaches to Semantically-Mediated Data Fusion.* Paper presented at the Military Communications Conference (MILCOM), Orlando, Florida, USA.
34. Haarslev, V., & Möller, R. (2003). Racer: A Core Inference Engine for the Semantic Web, *2nd International Workshop on Evaluation of Ontology-Based Tools*. Florida, USA.
35. Holger, K., Ferguson, R. W., Noy, N. F., & Musen, M. A. (2004) *The Protege OWL Plugin: An Open Development Environment for Semantic Web Applications.* Paper presented at the 3rd International Semantic Web Conference - ISWC 2004, Hiroshima, Japan.
36. Horrocks, I., Patel-Schneider, P. F., Boley, H., Tabet, S., Benjamin, G., & Dean, M. (2004). SWRL: A Semantic Web Rule Language Combining OWL and RuleML. World Wide Web Consortium: http://www.w3.org/Submission/SWRL/
37. Hustadt, U., Motik, B., & Sattler, U. (2005). Data complexity of reasoning in very expressive description logics, *19th International Joint Conferenceijcai on Artificial Intelligence (IJCAI 2005)*. Edinburgh, Scotland.

38. Kalfoglou, Y., Hu, B., Reynolds, D., & Shadbolt, N. R. (2005) *Semantic Integration Technologies Survey.* Southampton, England: Electronics and Computer Science, University of Southampton. (ECS e-Prints #10842)
39. Kifer, M. (2005). Requirements for an expressive rule language on the semantic web, *W3C Workshop on Rule Languages for Interoperability*. Washington D. C., USA.
40. Klir, G. J., & Yuan, B. (1994) *Fuzzy Sets and Fuzzy Logic: Theory and Applications*. Upper Saddle River, NJ, USA: Prentice-Hall Inc.
41. Kokar, M., Matheus, C., Baclawski, K., Letkowski, J., Hinman, M., & Salerno, J. (2004) *Use Cases for Ontologies in Information Fusion.* Paper presented at the 7th International Conference on Information Fusion, Stockholm, Sweden.
42. Lacy, L., Aviles, G., Fraser, K., Gerber, W., Mulvehill, A., & Gaskill, R. (2005). Experiences using OWL in military applications, *1st International Workshop on OWL Experiences and Directions*. Galway, Ireland.
43. Lacy, L., & Gerber, W. (2004). Potential modeling and simulation applications of the web ontology language-OWL, *36th Winter Simulation Conference (WSC 04)*. Washington D. C, USA.
44. Laera, L., Tamma, V., Bench-Capon, T., & Semeraro, G. (2004). SweetProlog: A System to Integrate Ontologies and Rules, *3rd International Workshop on Rules And Rule Markup Languages For The Semantic Web (RuleML2004)*. Hiroshima, Japan.
45. Llinas, J., Bowman, C., Rogova, G., Steinberg, A., Waltz, E., & White, F. (2004) *Revisiting the JDL Data Fusion Model II.* Paper presented at the 7th International Conference on Information Fusion, Stockholm, Sweden.
46. Matheus, C. (2005) *Using Ontology-based Rules for Situation Awareness and Information Fusion.* Paper presented at the W3C Workshop on Rule Languages for Interoperability, Washington D.C., USA.
47. Matheus, C., Kokar, M., & Baclawski, K. (2003a) *A Core Ontology for Situation Awareness.* Paper presented at the 6th International Conference on Information Fusion, Cairns, Queensland, Australia.
48. Matheus, C., Kokar, M., Baclawski, K., Letkowski, J., & Call, C. (2004) *The Practical Application of Semantic Web Technologies for Situation Awareness.* Paper presented at the International Semantic Web Conference, Hiroshima, Japan.
49. Matheus, C., Kokar, M., Baclawski, K., Letkowski, J., Call, C., Hinman, M., & Salerno, J. (2005a) *Lessons Learned From Developing SAWA: A Situation Awareness Assistant.* Paper presented at the 8th International Conference on Information Fusion, Philadelphia, USA.
50. Matheus, C., Kokar, M., Baclawski, K., Letkowski, J., Call, C., Hinman, M., Salerno, J., & Boulware, D. (2005b) *SAWA: An Assistant for Higher-Level Fusion and Situation Awareness.* Paper presented at the SPIE Conference on Multisensor, Multisource Information Fusion, Orlando, Florida, USA.
51. Matheus, C. J., Baclawski, K. P., & Kokar, M. (2003b) *Derivation of ontological relations using formal methods in a situation awareness scenario.* Paper presented at the Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications, Orlando, Florida, USA.
52. McKenzie, C., Preece, A., & Gray, P. (2006). Semantic Web Reasoning using a Blackboard System, *3rd International Workshop on Principles and Practice of Semantic Web Reasoning (PPSWR 2005)*. Dagstuhl Castle, Germany.
53. Meditskos, G., & Bassiliades, N. (2005). Towards an Object-Oriented Reasoning System for OWL, *1st International Workshop on OWL Experiences and Directions*. Galway, Ireland.
54. Meditskos, G., & Bassiliades, N. (2006). O-DEVICE: An Object-Oriented Knowledge Base System for OWL Ontologies, *4th Hellenic Conference on Artificial Intelligence* (Vol. 3955, pp. 256-266). Crete, Greece: Springer-Verlag.

55. Meditskos, G., & Bassiliades, N. (in press) A Rule-based Object-Oriented OWL Reasoner. *IEEE Transactions on Knowledge and Data Engineering*.

56. Mott, D., & Hendler, J. (2007) *Progress on the Collaborative Planning Model.* Paper presented at the 1st Annual Conference of the International Technology Alliance (ACITA), Maryland, USA.

57. Noy, N. F. (2004) Semantic integration: A survey of ontology-based approaches. *Sigmod Record, 33*(4), 65-70.

58. Noy, N. F., Sintek, M., Decker, S., Crubezy, M., Ferguson, R. W., & Musen, M. A. (2001) Creating Semantic Web Contents with Protege-2000. *IEEE Intelligent Systems, 16*(2), 60-71.

59. O'Connor, M., Tu, S., Nyulas, C., Das, A., & Musen, M. (2007). Querying the Semantic Web with SWRL, *International RuleML Symposium on Rule Interchange and Applications (RuleML 2007)*. Orlando, Florida.

60. O'Connor, M. J., Knublauch, H., Tu, S. W., Grossof, B., Dean, M., Grosso, W. E., & Musen, M. A. (2005) *Supporting Rule System Interoperability on the Semantic Web with SWRL.* Paper presented at the 4th International Semantic Web Conference (ISWC2005), Galway, Ireland.

61. Pan, J. Z., Stamou, G., Tzouvaras, V., & Horrocks, I. (2005). f-SWRL: A Fuzzy Extension of SWRL, *International Conference on Artificial Neural Networks (ICANN 2005)*. Warsaw, Poland.

62. Papamarkos, G., Poulovassilis, A., & Wood, P. T. (2003). Event-Condition-Action Rule Languages for the Semantic Web, *1st International Workshop on Semantic Web and Databases*. Berlin, Germany.

63. Parsia, B., & Sirin, E. (2004). Pellet: An OWL DL Reasoner, *International Workshop on Description Logics*. British Columbia, Canada.

64. Polleres, A. (2007). From SPARQL to rules (and back), *Proceedings of the 16th international conference on World Wide Web*. Alberta, Canada.

65. Preece, A., Gomez, M., de Mel, G., Vasconcelos, W., Sleeman, D., Colley, S., & La Porta, T. (2007) *An Ontology-Based Approach to Sensor-Mission Assignment.* Paper presented at the 1st Annual Conference of the International Technology Alliance (ACITA), Maryland, USA.

66. Pulvermacher, M., Semy, S., Obrst, L., & Stoutenburg, S. (2005) *Perspectives on Applying Semantic Web Technologies in Military Domains.* Bedford, Massachusetts, USA: The MITRE Corporation.

67. Richbourg, R., & Olson, W. K. (1996) A hybrid expert system that combines technologies to address the problem of military terrain analysis. *Expert Systems with Applications, 11*(2), 207-225.

68. Sattler, U. (2007). Reasoning in Description Logics: Basics, Extensions, and Relatives, *Reasoning Web*. Dresden, Germany.

69. Scherl, R., & Ulery, D. (2004) *Technologies for army knowledge fusion.* US Army Research Laboratory. (ARL-TR-3279)

70. Schreiber, G., Akkermans, H., Anjewierden, A., de Hoog, R., Shadbolt, N. R., Van de Velde, W., & Weilinga, B. (2000) *Knowledge Engineering and Management: The CommonKADS Methodology*. Massachusetts: MIT Press.

71. Semy, S. K., Pulvermacher, M. K., & Obrst, L. J. (2004) *Toward the Use of an Upper Ontology for US Government and US Military Domains: An Evaluation.* Bedford, Massachusetts, USA: The MITRE Corporation.

72. Serre, T., Wolf, L., Bileschi, S., Riesenhuber, M., & Poggio, T. (2007) Robust object recognition with cortex-like mechanisms. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 29*(3), 411-426.

73. Shadbolt, N., Hall, W., & Berners-Lee, T. (2006) The semantic web revisited. *IEEE Intelligent Systems, 21*(3), 96-101.

74. Shadbolt, N. R., & Burton, M. (1990) Knowledge elicitation: a systematic approach. In J. R. Wilson & E. N. Corlett (Eds.), *Evaluation of Human Work: A Practical Ergonomics Methodology* (2nd ed., pp. 406-440). London, England: Taylor and Francis.
75. Shadbolt, N. R., Gibbins, N., Glaser, H., Harris, S., & schraefel, m. c. (2004) CS AKTive Space: or how we learned to stop worrying and love the Semantic Web. *IEEE Intelligent Systems, 19*(3), 41-47.
76. 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.
77. Shadbolt, N. R., Tennison, J., Milton, N., & Howells, H. (2000) Situation Assessor Support System: A Knowledge-Based Systems Approach to Pilot Aiding. In P. T. McCabe, M. A. Hanson & S. A. Robertson (Eds.), *Contemporary Ergonomics*. London, UK: Taylor & Francis.
78. Sirin, E., Parsia, B., Grau, B. C., Kalyanpur, A., & Katz, Y. (2007) Pellet: A practical OWL-DL reasoner. *Web Semantics: Science, Services and Agents on the World Wide Web, 5*(2), 51-53.
79. Smart, P. R. (2001) *Development of a Tactical Decision Support System for FOAEW: AEW Barrier Positioning.* Nottingham, England: Epistemics.
80. Smart, P. R. (2002) *Development of a Tactical Decision Support System for FOAEW: Mission Planning.* Nottingham, England: Epistemics.
81. Smart, P. R. (2005) *Knowledge-Intensive Fusion for Improved Situational Awareness: Requirements for Knowledge-Filtered Awareness.* Southampton, England: School of Electronics and Computer Science, University of Southampton. (DTC/WP270/Requirements)
82. Smart, P. R., Bahrami, A., Braines, D., McRae-Spencer, D., Yuan, J., & Shadbolt, N. R. (2007a) *Semantic Technologies and Enhanced Situation Awareness.* Paper presented at the 1st Annual Conference of the International Technology Alliance (ACITA), Maryland, USA.
83. Smart, P. R., Russell, A., Shadbolt, N., schraefel, m. c., & Carr, L. A. (2007b) AKTiveSA: A Technical Demonstrator System for Enhanced Situation Awareness in Military Operations Other Than War. *The Computer Journal, 50*(6), 703-716.
84. Smart, P. R., & Shadbolt, N. R. (2007) *The Semantic Battlespace Infosphere: A Knowledge Infrastructure for Improved Coalition Inter-operability.* Paper presented at the 4th International Conference on Knowledge Systems for Coalition Operations (KSCO), Massachusetts, USA.
85. Smart, P. R., Shadbolt, N. R., Carr, L. A., & schraefel, m. c. (2005) *Knowledge-Based Information Fusion for Improved Situational Awareness.* Paper presented at the 8th International Conference on Information Fusion, Philadelphia, USA.
86. Steinberg, A., & Bowman, C. (2004) *Rethinking the JDL data fusion levels.* Paper presented at the National Symposium on Sensor and Data Fusion.
87. Stoilos, G., Simou, N., Stamou, G., & Kollias, S. (2006) Uncertainty and the Semantic Web. *IEEE Intelligent Systems, 21*(5), 84-87.
88. Stoilos, G., Stamou, G., Tzouvaras, V., Pan, J. Z., & Horrocks, I. (2005). Fuzzy OWL: Uncertainty and the Semantic Web, *International Workshop of OWL: Experiences and Directions*. Galway, Ireland.
89. Stoutenburg, S., Obrst, L., Nichols, D., Peterson, J., & Johnson, A. (2005). Toward a Standard Rule Language for Semantic Integration of the DoD Enterprise, *W3C Workshop on Rule Languages for Interoperability* (pp. 27-28). Washington D. C., USA.
90. Stuckenschmidt, H., & Klein, M. (2004). Structure-based partitioning of large concept hierarchies, *3rd International Semantic Web Conference (ISWC 2004)*. Hiroshima, Japan.
91. Sutherland, S. (1994) *Irrationality: Why We Don't Think Straight!* New Brunswick, NJ: Rutgers University Press.
92. Sycara, K., Paulucci, M., & Lewis, M. (2003) *Information discovery and fusion: semantics on the battlefield.* Paper presented at the 6th International Conference on Information Fusion, Cairns, Queensland, Australia.

93. Tsarkov, D., & Horrocks, I. (2006). FaCT++ description logic reasoner: System description, *International Joint Conference on Automated Reasoning (IJCAR 2006)*. Seattle, USA.
94. *The UK Joint High Level Operational Concept: An Analysis of the Components of the UK Defence Capability Framework.* (2005) Joint Doctrine and Concepts Centre, Ministry of Defence, UK. (Joint HLOC)
95. Wason, P. C. (1960) On the failure to eliminate hypotheses on a conceptual task. *Quarterly Journal of Experimental Psychology, 12*, 129-140.

Appendix A Acronyms & Abbreviations

AEW	Airborne Early Warning
AI	Artificial Intelligence
C2	Command and Control
COOL	CLIPS Object-Oriented Language
CWA	Closed World Assumption
CWM	Closed World Machine
CWR	Closed World Reasoning
DL	Description Logic
DLTR	Direct Loading of Triples and Rules
DSS	Decision Support System
DTD	Document Type Definition
ECA	Event Condition Action
FCG	Foreign Clearance Guide
FiRE	Fuzzy Reasoning Engine
f-OWL	Fuzzy OWL
f-SWRL	Fuzzy SWRL
IBM	International Business Machines Corporation
IEDM	Information Exchange Data Model
ILR	Incremental Loading of Rules
ILT	Incremental Loading of Triples
ITA	International Technology Alliance

JDBC	Java Database Connectivity
JDL	Joint Directors of Laboratories
JEOD	Joint Explosive Ordnance Disposal
Jess	Java Expert System Shell
LCWA	Local Closed World Assumption
M2M	Machine-To-Machine
MOOTW	Military Operations Other Than War
NATO	North Atlantic Treaty Organization
NEC	Network Enabled Capability
OWA	Open World Assumption
OWL	Web Ontology Language
OWR	Open World Reasoning
RDF	Resource Description Framework
RDFS	RDF Vocabulary Description Language
RIF	Rule Interchange Format
RuleML	Rule Markup Language
SPARQL	Simple Protocol and RDF Query Language
SQL	Structured Query Language
SWAP	Semantic Web Application Platform
SWRL	Semantic Web Rule Language
TTP	Tactics, Techniques and Procedures
USAF	United States Air Force

W3C	World Wide Web Consortium
WWW	World Wide Web
XML	eXtensible Markup Language
XMP	eXtensible Metadata Platform
XSLT	eXtensible Stylesheet Language Transformations