

O

T

S

D

AR-010-397

Conceptual Reasoning and Defence
Applications

Peter Deer, Peter W. Eklund &
Chris Nowak

DSTO-RR-0118

APPROVED FOR PUBLIC RELEASE

© Commonwealth of Australia

DTIC QUALITY INSPECTED 8

DEPARTMENT OF DEFENCE
DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

19980430 153

Conceptual Reasoning and Defence Applications

Peter Deer¹, Peter W. Eklund² and Chris Nowak³

¹Information Technology Division
Electronics and Surveillance Research Laboratory

²The Department of Computer Science
The University of Adelaide

³The School of Mathematics, Physics, Computing and Electronics
Macquarie University

DSTO-RR-0118

ABSTRACT

This report presents details of a formal technique called *conceptual reasoning*. Conceptual reasoning is a knowledge representation and reasoning framework for multiple-agent belief revision. Because conceptual reasoning has a particular interest in dealing with *contradictory* and *partial* information, it lends itself to applications in Defence Intelligence. We demonstrate the practical application of conceptual reasoning in a particular Defence Intelligence domain: multi-sensor fusion.

APPROVED FOR PUBLIC RELEASE

DTIC QUALITY INSPECTED 3

D E P A R T M E N T O F D E F E N C E

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

DSTO-RR-0118

Published by

*DSTO Electronics and Surveillance Research Laboratory
PO Box 1500
Salisbury South Australia 5108*

*Telephone: (08) 8259 5555
Fax: (08) 8259 6567
© Commonwealth of Australia 1997
AR No. AR-010-397
December 1997*

Approved for public release

Executive Summary

This report presents details of a formal technique called *conceptual reasoning*. Conceptual reasoning is a knowledge representation and reasoning framework for multiple-agent belief revision. Because conceptual reasoning has a particular interest in dealing with *contradictory* and *partial* information, it lends itself to applications in defence intelligence. We demonstrate the practical application of conceptual reasoning in a particular defence intelligence domain: multi-sensor fusion. Other applications foreseen for Defence include the analysis of: resource inventory, command and control structures, and command and control vocabulary.

We present the theory of *formal concept analysis* (FCA). A *context* in FCA is defined as consisting of a set of *objects* G , a set of *attributes* M , and an *incidence function* I which assigns attributes to objects by mapping $G \times M$ to the values $\{0, 2\}$. In FCA a related notion is that of a *concept*; given a context K its set of concepts $\underline{\mathcal{L}}(K)$ forms a *complete lattice*, called a *concept lattice* of K . We also consider *partial contexts* and *abstract contexts*, involving *abstract objects*. Then *valid sentences/descriptions* say which abstract objects are present in the context, and which are not. Given a set of descriptions $D_i \subseteq D$, the set D_i determines a *formal system* with *axioms* D_i and inference rules Φ . Hence, we introduce formal systems, and a set of their theories \mathbb{T} is equipped with an information ordering relation \leq .

Theories (T_i), when associated with *agents*, are called *believed theories* \mathbb{B} , and there is a minimal lattice \mathbb{C} of theories that includes the believed theories, but also contains *meets* and *joins* of the believed theories, $\mathbb{C} = \text{Cl}_{\wedge, \vee}(\mathbb{B})$. Such lattices provide a framework for common sense knowledge representation and reasoning. In particular, given such a lattice \mathbb{C} , it captures *contradictions* and *partiality*; from \mathbb{C} we can consider $\mathbb{C}_+ = \mathbb{C} \cup \{0, 1\}$. If $T_2 \geq T_1$ then T_1 is more *partial* than T_2 , and if $T_2 \vee T_1 = 1$ then the theories are contradictory. We also have that if, the lattice \mathbb{C}_+ is a *concept lattice* then, theories in \mathbb{C} and sentences in D can be *partially ordered*, and a *numerical measure* $\nu : \mathbb{T} \rightarrow [0, 1]$ can be derived. Hence, given that theories can be ordered (partially ordered, linearly ordered by their numeric measure), we can derive *preference relations* on theories which allows us to decide which information to accept or reject, or, more generally, how to order information, which in our case is expressed by theories.

The theoretical foundations of Conceptual Reasoning are presented. We propose *partial abstract contexts* as partial models of worlds consisting of objects having attributes. Further, we defined formal systems which allow us to find sentences (descriptions) that follow from a given set, and we relate the two, taking into account that the set of contexts \mathbb{K} and the set of theories \mathbb{T} are sets equipped with their information orderings.

The framework allows us to represent and reason about *contexts* involving *objects* having *attributes*. Both *semantic* and *syntactic* information about such contexts is considered. *Models* correspond to contexts, and syntactic information gives rise to *theories* consisting of sentences *provable* in *formal systems*. We then consider the multiple agent case. Given a set of agents, the set of their *believed theories* is \mathbb{B} , and the resulting lattice is \mathbb{C} , or \mathbb{C}_+ , if we include an empty theory 0 and an "inconsistent theory" 1 . We show how one should interpret such structures, and suggest how preference relations on theories can be derived.

We then focus on theories and consider related structures which allow to deal with information provided by *multiple agents*. As an illustration of the formalism, we show how conceptual reasoning can be applied in the multi-sensor fusion domain. The complete working of the example is provided as an appendix to this report.

Authors

Peter Deer

Information Technology Division

Peter Deer is a Senior Professional Officer in Information Technology Division. He joined DSTO in 1989 after 16 years in the Australian Regular Army. He pursues research interests in logic and human reasoning, and digital image processing using remotely sensed data.

Peter W. Eklund

The Department of Computer Science
The University of Adelaide

Peter Eklund is a Senior Lecturer in Computer Science at the University of Adelaide, where he teaches Artificial Intelligence and Logic Programming. He has recently been appointed as Foundation Professor of Information Technology at Griffith University.

Chris Nowak

The School of Mathematics, Physics, Computing
and Electronics
Macquarie University

Chris Nowak is a Research Associate at the School of Mathematics, Physics, Computing and Electronics at Macquarie University. His research interests are AI logic and mathematical algebra. He has recently completed a PhD in Computer Science at the University of Adelaide.

Contents

1	INTRODUCTION	1
1.1	Structure of the Report	1
1.2	Methodology	1
1.3	Multiple Agencies	1
1.4	Information Retrieval	3
1.5	Mathematical Foundations	4
1.6	Data Fusion	4
1.7	A Study of the Organisational Structure	6
2	CONCEPTUAL REASONING: BACKGROUND	7
2.1	Formal Concept Analysis	7
2.2	Fundamental Theorem of Concept Lattices	8
3	CONCEPTUAL REASONING: THEORETICAL FOUNDATIONS	11
3.1	Partial Abstract Contexts	11
3.2	Formal Theories	13
3.3	Contexts and theories	15
3.4	Multiple agents	16
4	INTELLIGENCE ANALYSIS AND CONCEPTUAL REASONING	21
4.1	Example: Multi-sensor Fusion	21
4.2	Incorporation of Database Knowledge	24
5	CONCLUSION	27
	References	28

1 INTRODUCTION

1.1 Structure of the Report

Section 1 discusses the broad approach of *Conceptual Reasoning* and its applications to the defence science domain. Formal concept analysis is introduced in Section 2.1 as the necessary background for the development of conceptual reasoning. The report develops the theory of Conceptual Reasoning in Section 3. A proof-theoretic account is offered in Section 3.2. Contexts are related to theories in Section 3.3. Multiple agency is addressed in Section 3.4 and an example follows in Section 4 in the multi-sensor fusion domain. Issues regarding the pragmatics of the technique are also examined in this section. The fully worked solution to the example presented in Section 4 is offered in an appendix.

Many standard techniques in engineering have been derived from early work in Artificial Intelligence (AI). Neural classifiers [HP91], fuzzy controllers [DH94], genetic algorithms [Hol75], classifier systems [Gol89] and knowledge and data discovery techniques [FU95] all have evolved from early AI research but are now mainstream engineering. The core AI literature is increasingly distancing itself from pragmatic tools, concentrating on theoretical features of disembodied intelligence [IJC95, AAA96, KR96] and commonsense reasoning. This ignores a pragmatic agenda, a convergence of computational intelligence and AI as a recognised engineering discipline. Such a convergence, if it were to be achieved, could provide a useful grab-bag of problem solving techniques for intelligent systems. We consider this goal highly desirable.

The aim of this paper is to present, in formal terms, the foundations of conceptual reasoning. Additionally, we give illustrations of potential applications of value in the defence science domain.

1.2 Methodology

Conceptual reasoning [Now97] described in this report is not mainstream AI, but it borrows heavily from mainstream AI methodology [GN87]. It deals with default or partial reasoning and multiple agencies, and uses logical foundations commonly used in AI logic research. It constructs a syntactic and semantic interpretation of multiple agencies (or worlds) unified via a soundness and completeness result.

1.3 Multiple Agencies

What is novel is the structure and ordering of implicit knowledge about the world according to order-theoretic principles. For example, since signal or sensor input often occurs as phenomena paired with several features, a useful approach to structuring information sources is to deal with *formal concepts*, discussed in the background section on page 7.

These formal concepts represent a cluster of attributes in which it would serve no useful purpose to distinguish. For example, in the UK, a precursor to a green light is a

flashing amber signal. Why is a warning signal necessary for a motorist starting from rest to change to a moving position? Unless construed as a wake up call to the driver, the amber elicits no response or action. It therefore has no information content. The amber is paired implicitly with the impending green light and signals the driver to motion. The example is anecdotal but it serves our purpose.

Another example derives from Australian newspapers. Since Australian newspapers tend to be controlled and run by only a few individuals, there is some content repetition in titles. Take *The Australian* and *The Adelaide Advertiser*, they often run identical stories and even images. If we consider each newspaper article as an agent, then both agents may report identically. The relative information content of the second is zero so we cluster both articles and their sources to the same source since they represent the same world view.

The collapse of multiple agreeing agencies may only be important in terms of the cardinality of the agreeing sources. If 10 sources report an incident X and 2 report something other than X (or nothing at all), this may represent the formation of a cartel for the belief of X . If cardinality can be interpreted as evidence in favour of X , there may be overwhelming evidence for X over something other than X .

A more interesting situation is where multiple sources indicate immediate or partial conflict. Where agents agree or disagree can be visualised as a partial order and this order represents both the inter-dependence of causal assumptions as well as the universal or agreed truths between agents. With conceptual reasoning, both contradictions, as well as agreement, can be isolated and graphically rendered. This rendering represents one of the key benefits of conceptual reasoning.

In terms of newsfeeds, many of the formal techniques in conceptual reasoning rely on a suitable translation of natural language to an agreed inter-lingua represented in logical terms. Each new report represents a theory and these theories are ordered in the same way as the paired signals above. This is perfectly tractable but relies on an appropriate representation of the theories in an agreed language sourced from the newsfeeds. This is a difficult problem. Since there is no unified theory of intention in natural language [RN95], there are only two current possibilities for automated reasoning of this type.

1. The first is a coarse machine translation to a universal logic such as predicate calculus. However, automated machine translation is an even more difficult problem. Such a translation will not only be flawed in terms of its inability to detect idioms, nuances and intentions (especially in deceptive sources) but worse still, since the predicate calculus is semi-decidable, we can never guarantee that a non-theorem does not follow.
2. Secondly, we could hand-craft newsfeeds into a tractable logical language, that should have the properties of soundness and completeness, say a description logic based on Knowledge Interchange Format (KIF) [GFB⁺92] or (a subset) of conceptual graphs [Sow83]. This would be an impossible task for the human analyst even for the most impoverished provincial newspapers. It would also lack the presentational adequacy of the original source, not everything, therefore, could be faithfully transcribed in the logical language and some of the information content would be lost [Gin93].

The idea of allowing the analyst to perform knowledge markup has been used by Philippe Martin and others [Mar95, MA96, Mar96]. Martin uses a knowledge markup language, based on conceptual structures, that allows the user to embed representation sentences within text documents. These sentences can then be used to navigate text or other documents elements, that is they can be used to hyperlink the source. Additionally, the mark up language can be used to make new inferences. Say we have two legal terms, “dissolution of company partnership” and “divorce settlement”. These two ideas combine to form the conjunction (or join), “divorcing partners who also have a company partnership”. The join of the first two can be used to search for instances of the join in case law.

1.4 Information Retrieval

An alternative to both machine translation and knowledge markup is to concentrate on efficient mechanisms to retrieve information on the basis of content addressability.

This approach places greater emphasis on the human analyst who must draw the appropriate inferences from the information retrieval process. This emphasis can be minimised by providing an inference engine that the analyst uses to select suitable facts from news sources and placing them into a “sheet of assertion”. This sheet represents those facts extracted from the database which are meaningful to the analytic task. Relevant facts can be combined to draw conclusions. In this sense, inferences can be guaranteed to be sound but in no way complete. This approach presents the most promising and tractable paradigm to intelligence analysis.

Most information retrieval (IR) relies on keyword boolean search, vector space models or n -grams [Dam95] as a mechanism for measuring document similarity [WMB94]. Conceptual structures [Sow83], although not conceived as an aid for information retrieval can be used as such. Conceptual structures are one of a number of candidate knowledge representation languages that rely heavily on order sorted relations to describe key concepts and relation types. A taxonomy of terms, as concepts (such as *animate-object* > *vehicle* > *car* > *BMW* > *BMW#123*) or connecting relations (such as *relationship* > *parent* > *mother* > *single-mother* > *single-mother-one-child* > *single-mother-one-child(mary, tom)*) can be defined and these taxonomies aid in the formation of partial orders of sentences. For example, when sufficiently sanitised, sentences in natural language can be visualised as a partial order or taxonomy of sorted data from unstructured text input. This ignores the semantic meaning of the sentences but orders them according to their interconnected relations and concept types.

The importance of this idea is that it allows information retrieval (IR) to be based not just on boolean keyword or vector space similarity, but also on the structural content of the information itself; and, in so far as semantics can be determined from structure, through the use of taxonomies of relations and concept types, meaning can be extracted. Furthermore, partial matches on specific information queries can be recalled. Precision of the information request is also thought to be improved since recall is based on either a specialisation or generalisation of the information request. This approach to IR is called

order-sorted or *knowledge retrieval* and is closely related to the literature on terminological logics [BFH⁺94] (sometimes called description logics).

In the United States, a company called Textwise (www.textwise.com) developed the first search engine to capitalise on this idea. Their search engine, called DR LINK, is now being used in the US Patent office and DoD agencies for IR against technical abstracts and other unstructured data. However, DR LINK is based on a carefully engineered taxonomy of parts of speech [LJ91], a so called sub-grammar ontology. Therefore its precision, in classical IR terms, is relatively poor.

Other companies such as the CYC Corporation (www.cyc.com) have been developing more general taxonomies of concepts and relations and companies such as Infoseek (www.infoseek.com) and Yahoo (www.yahoo.com) have developed Internet search engines based on ontologies that partition the Internet into various corpora of subject matter. There are several ANSI standard organisations addressing the issue of generalised term taxonomies [Pe93]. Apart from CYC, none we know suggest synthesising term hierarchies from the data sources themselves, i.e. via knowledge and data discovery techniques. A key premise reinforcing our work is that unstructured text sources can themselves be the sources for term hierarchies necessary for knowledge retrieval.

1.5 Mathematical Foundations

Conceptual reasoning, presented in this paper formally, combines and extends the work of two Mathematics Professors, Peter Burmeister [Bur89] and Bernhard Ganter [Gan96] from the Darmstadt group in Germany. It demonstrates how a natural preference ordering can be synthesised on propositional statements from various agents. The resulting formal framework provides, among other things, the capacity to visualise the process of assessing shared beliefs about the world, the so called common-knowledge that agents share. Additionally, it can isolate knowledge between agents that is contradictory and will recommend the cleanest way to resolve such conflicts through retracting agent beliefs. Nowak's work [Now97] represents the extreme end of the continuum between applications of formal concept analysis [Wil92], and computational intelligence. Conceptual reasoning is much closer to traditional symbolic AI but has a strong theoretical connection to general algebra and lattice and order theory that importantly, unlike some other belief revision and nonmonotonic reasoning frameworks, imposes a natural ordering on beliefs.

1.6 Data Fusion

In almost all respects, conceptual reasoning, for the task of processing and interpreting the implications of text-based information sources, is 5 years from the intelligence analyst's desktop. The problems of engineering a suitable knowledge retrieval framework for text-based sources is a more pressing pre-requisite for the success of the approach and will take considerable engineering and research effort. However, the conceptual reasoning framework may prove to be of value when presented with multiple and ambiguous signal processing inputs in the sensor and data fusion task or for the analysis of organisational structure.

Data fusion is the activity in which related data from multiple sensors or sources is combined to provide enhanced quality and availability of information over that which is available from any individual sensor or source. There have been various architectures proposed for data fusion, but the commonly held view is that the fusion process can be divided into a number of levels. For example, the modified US JDL Model breaks down the process of data fusion in the military context into four levels, namely object refinement, situation refinement, threat refinement and process refinement — see [WL90] and [NM96].

Object refinement is an iterative process of fusing data to determine the identity and other attributes of entities and also build tracks to represent their behaviour. The product from this level is called the *situation picture*. *Situation refinement* is an iterative process of fusing the spatial and temporal relationships between entities to group them together and form an abstracted interpretation of the distribution of forces. The product from this level is called *Situation Assessment*. *Threat Assessment* is an iterative process of fusing the combined activity and capability of enemy forces to infer their intentions and assess the threat they pose to Own and allied forces. The product from this level is called *Threat Assessment*. Finally, *Process Refinement* is an iterative process of improving the products from other fusion levels. The performance is monitored and when it degrades sufficiently the collection plans are adapted so that the necessary sensor and source data can be obtained. This is not a fusion process in itself, but is an integral part of any data fusion system. Results from a higher level can also be used to indirectly influence the fusion process in lower levels.

Situation assessment has meaning at the tactical, operational and strategic level. A complete understanding of the theatre over time may reveal the tactical intentions of the enemy but it may be impractical to undertake in real-time. Alternatively, *conceptual reasoning* could also be used as a representational approach in order to assess previously collected kinematic and attribute operational data and organising that data using unsupervised machine learning. In much the same way as this approach can be used to examine chess games to isolate tactical chunks, so too conceptual reasoning can be used to isolate, and later recognise, tactical threats.

Above this architecture for data fusion, our interest is in optimising resource utilisation. The analysis of historical theatre data can be used to determine if resources were used optimally or sub-optimally. A re-organisation of the resource allocation during operational manoeuvres will reveal the types and source of data required to perform real-time operational assessment through a static analysis of the historical theatre. Again, approaches similar (or identical) to conceptual reasoning may be appropriate to this task.

The key notion reinforcing these possibilities is an answer to the question: can a cross table or context be created from the input data sources? In data fusion, the creation of the formal context needs to be tied to the physics of the sensor device(s). Transitions between states may need explanation and this gives rise to considerable complexity between states of the sensor, their explanation and points in time.

One general, and incorrect assumption, concerning conceptual reasoning, and formal concept analysis more generally, is that it can not represent time points or intervals: it can. Contexts can be created for various timepoints or intervals and the inter-connection between these contexts examined [Wil96].

1.7 A Study of the Organisational Structure

Structure, and the visualisation of structure, is the main selling point of conceptual reasoning and formal concept analysis. The capacity to visualise and explain structure and intentions through lattice diagrams is the key to developing an understanding of an analytical domain. Such visualisation capabilities can be used to expose flaws and weaknesses in reasoning and organisation. There are three ways we see this idea being manifest: (i) resource inventory; (ii) lines of command and communication; (iii) the vocabulary for command and control. All relate to formal concept analysis in the first instance, but could, if so desired, be made more sophisticated to take into account the conceptual reasoning framework through the modelling of the object's internal goals or intentions as agent theories.

The first is relatively straightforward. It is easy to imagine a cross table (such as the one on page 8) containing all operational units and the resources (or capabilities) at their disposal. How those units cluster in the concept lattice tells us something about the command/capability relationship. Secondly, lines of command and structure can be built into a formal concept lattice, the natural and intended structure of the organisation should be revealed by that process. The process would not be expected to say much about a carefully organised command structure, but, if and when there are ambiguities, these can be isolated and resolved with little discussion. In this respect, a mechanised and formal analytical process such as formal concept analysis (or conceptual reasoning) has great advantages in defusing political debate. Dispute is minimised when organisation structure can be expressed as crosses in a matrix. The interpretation of the synthesised structure is, once again, unambiguous.

Finally, the study of the structure and organisation of command and control languages is an important element of the efficiency and battle-readiness of the military. Formal concept analysis provides a tool for such analysis. Frank Vogt [Vog97] has shown how formal concept analysis can be used as an aid in communication for the process of software engineering. He shows how "problem", "design" and "implementation" cross tables can be created representing multiple and different views of the engineering process. The effect of using formal concept analysis in this way is to reveal a structure that directly isolates the project priorities through task inter-dependence. Since different views of the operational unit's capabilities are similarly defined in the military domain, this analytic process will once again have the effect of identifying in precise terms the priorities of the operational unit.

The prospects for the application of conceptual reasoning in the military domain are excellent. However, before we can generate a capability with this technology, experience needs to be built using formal concept analysis as the starting point.

We now give the detailed formalism starting with an introduction to formal concept analysis.

2 CONCEPTUAL REASONING: BACKGROUND

2.1 Formal Concept Analysis

Formal Concept Analysis is a theory of concept formation derived from lattice and order theory. In Engineering and AI terms it is an unsupervised learning technique.

The principal proponents of Formal Concept Analysis (FCA) can be found in the *Fachbereich Mathematik* at the *Technische Hochschule Darmstadt* in Germany. FCA is inspired by Birkhoff's work in *Lattice Theory* from the 1930s and 40s and studies in Universal Algebra by Grätzer from the early 1970s. An excellent reading in the Introduction to Lattices and Order is given by Davey and Priestley [DP90]. Definitive papers on the subject of FCA by its inventor Rudolf Wille are [Wil82, Wil92, GW96].

FCA asks the question, what is a concept? One answer is that a concept is determined by its *intent* and *extent*. The extent describes all the objects in the universe that belong to a concept, e.g. the set of all "red Hondas", mine, yours, the guy's down the street, the one that just drove past. The intent is the collection of all attributes shared by a set of objects, e.g. the set of all red Hondas use fuel, they are all cars, they are all designed in Japan, presumably they all have radios, wheels, gears and so on.

Because a concept can have many instances, and the set of all instances is an almost limitless set of shared attributes of one sort or another, it is customary to work with a specific context in which both the set of objects (as ground instances) and attributes (or characteristics) are fixed. One of the profound research questions relevant to FCA is whether or not it can be used when contexts are "open" or, in AI terms, informationally incomplete. In simple terms, conceptual reasoning considers whether many contexts can be analysed to determine the extent to which they represent a coherent set of multiple-agent beliefs. Wille himself has considered this question [Wil96]. The difference between Nowak's and other work is the use of an abstract intent, which can allow for a countably infinite set of objects to be dealt with, and the use of multiple description sets. Nowak's approach is a model-theoretic AI approach which exploits the lattice-based "structure" of concept clusters but deals with them as members of an object language in logic.

The best way to understand FCA is to consider a simple example. In FCA, we always build a "cross table" or context in which all the objects and their attributes are enumerated. The following example can be found in Davey and Priestly's book on page 221.

In this example, objects are Planets in our solar system. Planets have attributes or attributes that are fixed in terms of their distance from the sun, their size and whether or not they have a moon(s). In table 1, the object in the i th row possesses the attribute in the j th column exactly when the (i, j) th cell is marked with an x. A concept is an ordered pair (A, B) where A is a subset of the 9 planets of the solar system and B is the set of all attributes shared by the objects in A . This also means that A is the set of all objects that possess each of the attributes in B . This means that the formal concepts denoted (A, B)

	SIZE			DISTANCE		MOON	
	small	medium	large	near	far	yes	no
<i>Mercury</i>	x			x			x
<i>Venus</i>	x			x			x
<i>Earth</i>	x			x		x	
<i>Mars</i>	x			x		x	
<i>Jupiter</i>			x		x	x	
<i>Saturn</i>			x		x	x	
<i>Uranus</i>		x			x	x	
<i>Neptune</i>		x			x	x	
<i>Pluto</i>	x				x	x	

Table 1: A cross table for the Solar system.

consist of just those objects with a set of attributes in B . For example, take the object *Earth* and consider its attributes;

$$B = \{small, near, yes\}.$$

Now ask, what are the other planets that possess all the attributes in B ? The answer is *Mars*. So one concept in this context is $\{\{Earth, Mars\}, \{small, near, yes\}\}$. The concept has as its extension of the set of planets $\{Earth, Mars\}$ and as its intention of the set of attributes $\{small, near, yes\}$.

Note the other concepts that result from this context:

$$\begin{aligned} &\{\{Mercury, Venus\}, \{small, near, no\}\}; \\ &\{\{Jupiter, Saturn\}, \{large, far, yes\}\}; \\ &\{\{Uranus, Neptune\}, \{medium, far, yes\}\}; \\ &\{\{Pluto\}, \{small, far, yes\}\}. \end{aligned}$$

Thus, although there are 7 attributes and 9 objects, there are only 5 distinct formal concepts.

It is also usual to regard a concept as more general than another if its extent is a subset of the other's concept extent. This means we can define an order on concepts $(A_1, B_1) \leq (A_2, B_2)$ iff A_1 is a subset of A_2 . This results in a partial ordering over concepts which turns out to have the attributes of a complete lattice.

2.2 Fundamental Theorem of Concept Lattices

A context is a triple (G, M, I) where G and M are disjoint sets and I is a subset of the Cartesian product of G and M , i.e. $G \times M$. The elements of G and M are *objects* and *attributes*. An object g has the property m , denoted gIm iff $(g, m) \in I$. The reason that

G and M are so named is that they are derived from the German for object (*Gegenstände*) and M for attribute (*Merkmale*).

For $A \subseteq G$ and $B \subseteq M$, define $A' = \{m \in M \mid \forall_{g \in A} gIm\}$ and $B' = \{g \in G \mid \forall_{m \in B} gIm\}$. Hence, A' is the set of all attributes common to the objects in A and B' is the set of all processing the properties in B . A formal concept is a pair (A, B) of the context (G, M, I) such that $A' = B$ and $B' = A$.

The set of all formal concepts of (G, M, I) is denoted by $\mathcal{B}(G, M, I)$. Again, following the German traditions of the formalism, the \mathcal{B} denotes the German "Begriff" which roughly translates as "description".

For concepts (A_1, B_1) and (A_2, B_2) in $\mathcal{B}(G, M, I)$ we write $(A_1, B_1) \leq (A_2, B_2)$ if $A_1 \subseteq A_2$. An important theoretical result is that this turns out to be equivalent to requiring $B_2 \supseteq B_1$. This implies that the structure $(\mathcal{B}(G, M, I), \leq)$ is a complete lattice, called the concept lattice of $\mathbf{B}(G, M, I)$.

Concepts are placed in a lattice structure in which the meet and join of any combination of elements are given by definition. This concept lattice not only contains concepts corresponding to each object but also concepts corresponding to the meet and join of other concepts. The lattice can express all relationships between objects and properties. For example, the lattice can represent the fact that if an object has one attribute, then it must possess all attributes lying above this node in the lattice. It is the capability to express such relationships that makes the lattice a powerful algebraic structure.

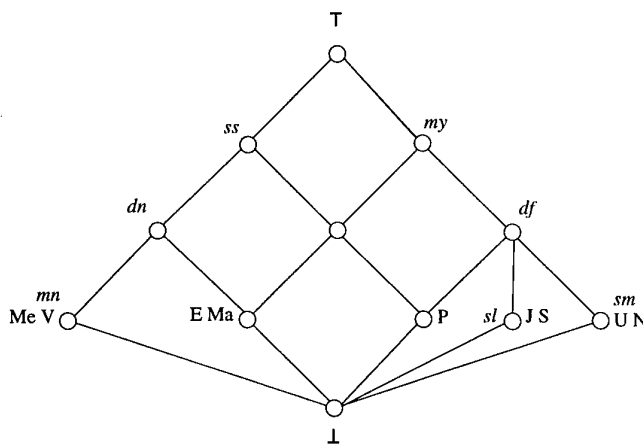


Figure 1: The formal concept lattice for Table 1.

For example, if we refer to the context described in the table above, we can look at the midpoint of the lattice and the fundamental theorem tells us that we can determine that this midpoint is $(\{Earth, Mars, Pluto\}, \{small, yes\})$.

The concept lattice provides a basic analysis of a context, it yields an appropriate classification of objects and at the same time indicates the implications between properties.

3 CONCEPTUAL REASONING: THEORETICAL FOUNDATIONS

3.1 Partial Abstract Contexts

A *context* in Formal Concept Analysis (FCA), see e.g., [DP90], is defined as a triple $K = (G, M, I)$, where G is a set of *objects*, M is a set of *attributes*, and I is an *incidence function* associating objects with attributes where $G \cap M = \emptyset$ and $I: G \times M \rightarrow \{0, 2\}$ is to be interpreted as follows.

$$I(g, m) = \begin{cases} 2 & \text{if } g \text{ has } m, \\ 0 & \text{if } g \text{ does not have } m. \end{cases}$$

Refer to such contexts as *standard FCA contexts*. It is natural to consider *partial*, or *three-valued* contexts, where the incidence function allows partially specified (or partially known) objects; in this case $I: G \times M \rightarrow \{0, 1, 2\}$ is to be interpreted as follows:

$$I(g, m) = \begin{cases} 2 & \text{if } g \text{ is known to have } m, \\ 0 & \text{if } g \text{ is known not to have } m, \\ 1 & \text{otherwise.} \end{cases}$$

Apart from the given set of attributes $P = \{p_1, \dots, p_n\}$, one might want to include the corresponding *negated attributes* denoted by $\bar{P} = \{\bar{p}_1, \dots, \bar{p}_n\}$. We assume that this is the case, so $M = P \cup \bar{P}$. It is then required that $I(g, m) = 2$ iff $I(g, \bar{m}) = 0$, and $I(g, m) = 1$ iff $I(g, \bar{m}) = 1$. We want to deal with such *partial contexts*.

However, it is often problematic to *uniquely identify* objects. Hence, given a set of attributes M , we employ a set of *abstract objects*—these can be associated with sets of *indiscernible* objects. First, we introduce a set \mathbf{F} of *formulae*, given as follows.

$$\mathbf{F} = \{F \subseteq M \mid \forall p \in P \ F \not\supseteq \{p, \bar{p}\}\}$$

Abstract objects can be associated with formulae. If $F \in \mathbf{F}$ then we define the *abstract object* \mathbf{g}_F to be the superset of G which has exactly the set F as its set of attributes. Hence, the set \mathbf{G} of all abstract objects.

$$\mathbf{G} = \{\mathbf{g}_F \mid F \in \mathbf{F}\}$$

If $\mathbf{g}_F \in \mathbf{G}$ then \mathbf{g}_F is (*properly*) *partial* iff $|F| < |P|$, otherwise it is *total*. Further, \mathbf{G} can be seen as a set equipped with an information ordering relation \leq given by $\mathbf{g}_{F_1} \leq \mathbf{g}_{F_2}$ iff $F_1 \subseteq F_2$.

Let $\mathbf{G}_i \subseteq \mathbf{G}$. \mathbf{G}_i determines a context $\mathbf{K}_i = (\mathbf{G}_i, M_i, I_i)$ – assume that M_i is determined by the set of attributes of the objects of \mathbf{G}_i , and note that the objects also specify the incidence relation I_i . Then \mathbf{K}_i is referred to as an *abstract context*. Note also that because \mathbf{G}_i can include *partial objects*, \mathbf{K}_i is *partial*. Let \mathbb{G} denote the powerset of \mathbf{G} , i.e., $\mathbb{G} =$

$\mathcal{P}(\mathbf{G})$. Then every element of \mathbf{G} determines an abstract context. The set \mathbb{K} of all abstract context (over the fixed M) is denoted by

$$\mathbb{K} = \{K_i = (\mathbf{G}_i, M_i, I_i) \mid \mathbf{G}_i \in \mathbf{G}\}.$$

An example of an abstract context is presented in Table 2, in a form of a “+, −, • table”—if an object is known to have m (known to not have m) then this is denoted by + sign (by − sign), otherwise the • sign is used. For instance, the object $g_{F_{201}}$ has the attributes $IsSedan$ and \overline{IsRed} . Abstract contexts can be used to represent information about a world, without uniquely identifying objects—the context of Table 2 represents a world in which there are “red non-Fords” and there are “non-red sedans,” and there are no other objects.

$g_{F_{120}}$	•	+	−
$g_{F_{201}}$	+	−	•
K	IsSedan	IsRed	IsFord

Table 2: Abstract context—example

The set \mathbb{K} can be equipped with an *information ordering* defined as follows. Assume a fixed M , and let K_1 and K_2 be abstract contexts with sets of objects \mathbf{G}_1 and \mathbf{G}_2 , respectively. Then $K_1 \leq K_2$ if and only if the following two conditions are satisfied:

1. $\forall g_2 \in \mathbf{G}_2 \exists g_1 \in \mathbf{G}_1 g_1 \leq g_2$
2. $\forall g_1 \in \mathbf{G}_1 \exists g_2 \in \mathbf{G}_2 g_1 \leq g_2$

This turns \mathbb{K} into a partially ordered set (\mathbb{K}, \leq) . Maximal elements of \mathbb{K} are called *total contexts* and are denoted by \mathbb{K}^\uparrow . An example of (\mathbb{K}, \leq) for $|P| = 1$ is presented in Figure 2—the figure shows sets of abstract objects, but these determine the corresponding contexts.

In the example of Figure 2, if an object is *in, outside, on border* of the circle, then it *has, does not have, is undetermined* wrt p_1 . There are four maximal elements, one representing an *empty* world—these give rise to total contexts, the context of the empty world having an empty set of objects. There are two minimal elements, one representing a nonempty but otherwise unspecified world, the other representing an empty world.

Given a context $K_i \in \mathbb{K}$ and the set of formulae F . A formula F is defined to be \oplus -valid iff;

$$K_i \models_{\oplus} F \text{ iff}_{\text{def}} \exists g \in \mathbf{G}_i \forall m \in F I_i(g, m) = 2$$

Likewise F is \ominus -valid iff;

$$K_i \models_{\ominus} F \text{ iff}_{\text{def}} \forall g \in \mathbf{G}_i \exists m \in F I_i(g, m) = 0$$

Then, define a set D of *descriptions*, or sentences.

$$D = F \times \{\oplus, \ominus\}$$

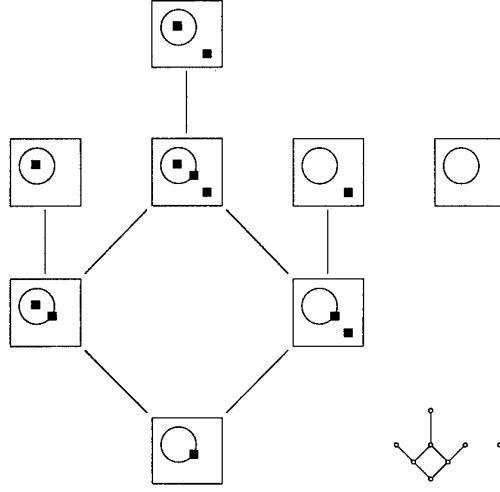


Figure 2: Information ordering on contexts

Descriptions are “marked formulae” D can be viewed as $\{\oplus F, \ominus F \mid F \in \mathbf{F}\}$. Validity of descriptions can then be defined in terms of \oplus - and \ominus -validity. We say that D is *valid* in \mathbf{K} , denoted $\mathbf{K} \models D$ iff one of the following conditions holds.

$$D = \oplus F \text{ and } \mathbf{K} \models_{\oplus} F, \text{ or}$$

$$D = \ominus F \text{ and } \mathbf{K} \models_{\ominus} F.$$

If \mathbf{K} is a context then we say that

$$\mathbf{T}_{\mathbf{K}} = \{D \in \mathbf{D} \mid \mathbf{K} \models D\}$$

is a *theory* of \mathbf{K} .

3.2 Formal Theories

Suppose that a set D_i of descriptions is given, i.e., $D_i \subseteq \mathbf{D}$. The intention is that D_i is a set of descriptions valid in some context. We propose *formal systems*: given a set D_i a formal system \mathcal{H}_i allows us to find descriptions that follow from D_i .

The following set $\Phi = \{\varphi_1, \varphi_2, \varphi_3, \varphi_4\}$ of *inference rules* is employed.

$$\varphi_1 : \frac{\oplus F \cup \{m\}}{\oplus F} \quad \varphi_2 : \frac{\ominus F}{\ominus F \cup \{m\}}$$

$$\varphi_3 : \frac{\ominus F \cup \{m\}, \ominus F \cup \{\bar{m}\}}{\ominus F} \quad \varphi_4 : \frac{\oplus F, \ominus F \cup \{m\}}{\oplus F \cup \{\bar{m}\}}$$

For instance, the rule φ_4 can take the form:

$$\frac{\oplus\{IsCar\}, \ominus\{IsCar, IsRed\}}{\oplus\{IsCar, IsRed\}}$$

which says that if there are “cars” and there are no “non-red cars” then there are “red cars.”

We remark that although different sets of description sets can be considered, the same set of inference rules Φ is employed. When defining a formal system \mathcal{H}_i , the set D_i is treated as a set of (proper) axioms, and the set of descriptions that follow from D_i using the inference rules Φ is called a *theory*, and is denoted by T_i , i.e., $T_i = Cn_{\Phi}(D_i)$ (the closure of D_i under $\Phi = \{\varphi_1, \varphi_2, \varphi_3, \varphi_4\}$ or simply $Cn(D_i)$). Hence, a description set $D_i \subseteq D$ is a theory if $D_i = Cn(D_i)$. A theory T_i is *consistent* if there is no $F \in F$ such that $T_i \supseteq \{\oplus F, \ominus F\}$. We say that a description set D_i is consistent if the theory $T_i = Cn(D_i)$ is. Given a theory T_i there is a minimal subset A_i of T_i such that $Cn(A_i) = T_i$; the set A_i is called a *minimal unique axiom set*, or a *generator* of T_i , and we write $A_i = gen(T_i)$.

An example of a description set D_i , its theory $T_i = Cn(D_i)$, and the generator $A_i = gen(T_i)$ of the theory is given below.

$$\begin{aligned} D_i &= \{ \oplus\{p_1\}, \ominus\{p_1, p_2, p_3\}, \ominus\{p_1, p_2, \bar{p}_3\} \} \\ T_i &= \{ \oplus\{p_1, \bar{p}_2\}, \oplus\{p_1\}, \oplus\{\bar{p}_2\}, \oplus\{\}, \ominus\{p_1, p_2\}, \\ &\quad \ominus\{p_1, p_2, p_3\}, \ominus\{p_1, p_2, \bar{p}_3\} \} \\ A_i &= \{ \oplus\{p_1, \bar{p}_2\}, \ominus\{p_1, p_2\} \} \end{aligned}$$

We limit ourselves to *consistent theories*: the set of all consistent theories (over a fixed M) is denoted by \mathbb{T} . That is,

$$\mathbb{T} = \{T_i \subseteq D \mid T_i \text{ is a consistent theory}\}$$

An obvious choice for an *information ordering* relation \leq on consistent theories is the relation of set inclusion—a bigger theory says more about the context. If $T_1, T_2 \in \mathbb{T}$ then

$$T_1 \leq T_2 \text{ iff } T_1 \subseteq T_2.$$

This turns \mathbb{T} into a *partially ordered set*, or *poset* (\mathbb{T}, \leq) . An empty set \emptyset of descriptions is denoted by $\mathbf{0}$. Clearly, $\mathbf{0}$ is a consistent theory, so $\mathbf{0} \in \mathbb{T}$. Maximal elements of \mathbb{T} are *total theories*, and are denoted by \mathbb{T}^\dagger .

Form the set $\mathbb{T}_1 = \mathbb{T} \cup \{\mathbf{1}\}$, where $\mathbf{1}$ is required to satisfy $T_i < \mathbf{1}$, for any $T_i \in \mathbb{T}$. Then we have the following,

Proposition 1 (\mathbb{T}_1, \leq) is a lattice.

Proof: It is a poset: the relation \leq is reflexive, antisymmetric and transitive, because \subseteq is. Supremum \vee and infimum \wedge are given by

$$T_1 \vee T_2 = \begin{cases} Cn(T_1 \cup T_2), & \text{if } T_1 \vee T_2 \text{ is consistent} \\ \mathbf{1} & \text{if } T_1 \vee T_2 \text{ is inconsistent} \end{cases}$$

$$T_1 \wedge T_2 = T_1 \cap T_2$$

The theory $\mathbf{1}$ can be seen as an element employed to *shift* \mathbb{T} to obtain \mathbb{T}_1 , but one can also think of $\mathbf{1}$ as representing inconsistent theories \square .

3.3 Contexts and theories

Recall that (\mathbb{K}, \leq) and (\mathbb{T}, \leq) are partially ordered sets. When relating contexts to theories we define mappings between them, consider whether the ordering is preserved, and comment on soundness and completeness.

Define a mapping $\tau : \mathbb{K} \rightarrow \mathbb{T}$ as follows. If $\mathbf{K}_i \in \mathbb{K}$ then $\tau(\mathbf{K}_i) = \mathbf{T}_{\mathbf{K}_i} = \{D \in \mathcal{D} \mid \mathbf{K}_i \models D\}$. Referring back to Section 3.1, where *theories of contexts* were defined, the mapping τ maps contexts to their theories. We easily obtain the following result, stating that theories of contexts are indeed consistent theories.

Proposition 2 *Let $\mathbf{K}_i \in \mathbb{K}$. Then $\mathbf{T}_{\mathbf{K}_i} = \tau(\mathbf{K}_i) \in \mathbb{T}$.*

Proof *The set of formulae valid in \mathbf{K}_i is closed under Φ , because any syntactic inferences are semantically valid - they preserve validity - hence, valid formulae of \mathbf{K}_i form a theory. The theory is consistent, because it is impossible for a context to have both $\oplus F$ and $\ominus F$ valid \square .*

An example showing contexts of Figure 2 mapped to theories is shown in Figure 3.

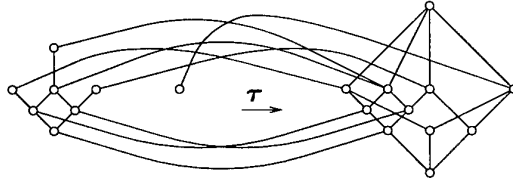


Figure 3: Contexts, theories and τ mapping

The right hand side of Figure 3 shows the lattice of all consistent theories, for the case $|P| = 1$. As convention the empty context—refer to Figure 2—is mapped to a theory that contains all descriptions of the form $\ominus F$, but none of the form $\oplus F$ —two theories covered by it (immediately below it) are not theories of any contexts, more precisely, they are $\{\ominus\{p_1\}\}$ and $\{\ominus\{\bar{p}_1\}\}$, and indeed, they provide some information about the context, without specifying whether or not the context is empty! Thus, the reverse mapping, from theories to contexts is more problematic—in fact we define a mapping from theories to *sets* of contexts.

Define a mapping $\kappa : \mathbb{T} \rightarrow \mathcal{P}(\mathbb{K})$ as follows. If $\mathbf{T}_i \in \mathbb{T}$ then $\kappa(\mathbf{T}_i)$ is given by $\kappa(\mathbf{T}_i) = \{\mathbf{K} \in \mathbb{K} \mid \mathbf{T}_{\mathbf{K}} \geq \mathbf{T}_i\}$, we use $\mathcal{K}_{\mathbf{T}_i}$ to denote $\kappa(\mathbf{T}_i)$. Further, if $\mathcal{D}_i \subseteq \mathcal{D}$ is consistent, and hence $\text{Cn}(\mathcal{D}_i) = \mathbf{T}_i \in \mathbb{T}$ then we can extend the mapping \mathbf{K} to all of \mathcal{D}_i by defining $\kappa(\mathcal{D}_i) = \kappa(\text{Cn}(\mathcal{D}_i)) = \kappa(\mathbf{T}_i)$

The set of *total models* for \mathbf{T}_i is denoted by $\text{MOD}(\mathbf{T}_i)$ and is given by $\text{MOD}(\mathbf{T}_i) = \mathbb{K}_i^\uparrow = \{\mathbf{K} \in \mathbb{K}^\uparrow \mid \mathbf{T}_{\mathbf{K}} \in \mathbb{T}_i^\uparrow\}$, where $\mathbb{T}_i^\uparrow = \{\mathbf{T} \in \mathbb{T}^\uparrow \mid \mathbf{T} \geq \mathbf{T}_i\}$. The set of *minimal partial models* for \mathbf{T}_i is denoted by $\text{mod}(\mathbf{T}_i)$, or $\mathcal{K}_{\mathbf{T}_i}$ and is given by $\text{mod}(\mathbf{T}_i) = \kappa(\mathbf{T}_i) = \{\mathbf{K} \in \mathbb{K} \mid \mathbf{T}_{\mathbf{K}}$

$\geq T_i$ and K is \leq -minimal}, and it is referred to as a κ -model of T_i . We say that D is *valid* in \mathcal{K}_{T_i} , denoted by $\mathcal{K}_{T_i} \models D$, if D is valid in every element of \mathcal{K}_{T_i} .

Regarding the information ordering, we easily obtain the following. Let $T_1 = \tau(K_1)$, $T_2 = \tau(K_2) \in \mathbb{T}$. Then we have that if $K_1 \leq K_2$ then $T_1 \leq T_2$. Indeed, getting additional information about the same context does not invalidate the descriptions that are already valid—this is the result of our definition of validity. One can easily notice that \leq cannot be replaced by its strict version $<$, as there are contexts such that one is strictly above the other, but both have the same theory: for instance, with respect to Figures 2 and 3, note that the only context of Figure 2 that has three objects and the context above it have the same theory.

The above considerations allow us to state a soundness and completeness theorem.

Proposition 3 *Let \mathcal{H}_i be a formal system with axioms D_i , let \mathcal{K}_i be a κ -model of D_i , and let $D \in \mathcal{D}$ and $D_i \subseteq \mathcal{D}_i$. Then:*

$$\mathcal{K}_i \models D \text{ iff } \mathcal{H}_i \vdash D$$

Proof (Sketch) *Soundness is immediate—if D_i is a set of axioms of \mathcal{H}_i and $\mathcal{H}_i \vdash D$ then it is sufficient to notice that the semantic equivalent of the syntactic proof can be carried out in every context of \mathcal{H}_i —hence, D is valid in \mathcal{H}_i . Completeness is more complicated, and employs a procedure for generating the κ -model for D_i \square .*

In the next section we limit ourselves to theories, i.e., we stay on the syntactic side of the formalism. Every theory however has its κ -model, a set of minimal contexts in which the theory is valid.

3.4 Multiple agents

Let S be a set of *agents* and let $\{D_s\}_{s \in S}$ be a set of description sets provided by the agents. Two agents might provide the same description set, and even if their description sets differ they might produce the same theory. Let \mathbb{B} be the set of believed theories of the agents—certainly, $\mathbb{B} \subseteq \mathbb{T}$. Assume believed theories are *nonempty* and *consistent* (agents are assumed to be consistent), so $\mathbb{B} \cap \{\mathbf{0}, \mathbf{1}\} = \emptyset$. We define $\mathbb{C} = \text{Cl}_{\wedge, \vee}(\mathbb{B})$, where \wedge and \vee are the operations already defined, Section 3.2. Further, we define $\mathbb{C}_+ = \mathbb{C} \cup \{\mathbf{0}, \mathbf{1}\}$, and so forming \mathbb{C}_+ simply accounts to adding—unless they are already there—a *bottom* $\mathbf{0}$ and a *top element* $\mathbf{1}$ to \mathbb{C} unless they are already there. It follows that \mathbb{C} and \mathbb{C}_+ are *lattices*, because they are subsets of \mathbb{T} closed under \wedge and \vee .

Consider an example. Let $\mathbb{B} = \{B_1, B_2\}$, where $B_1 = \text{Cn}(\{\oplus\{p_1, p_2\}\})$ and $B_2 = \text{Cn}(\{\ominus\{p_2\}, \oplus\{p_1, \overline{p_2}\}\})$. We get $B_1 \wedge B_2 = C_3 = \text{Cn}(\{\oplus\{p_1\}\})$ and $B_1 \vee B_2 = \mathbf{1}$ —the theories B_1, B_2 have a nonempty *meet* (intersection), but they cannot be *joined* consistently—see Figure 4(a). The closure \mathbb{C} of \mathbb{B} , i.e., $\mathbb{C} = \text{Cl}_{\wedge, \vee}(\mathbb{B})$ is the set $\{B_1, B_2, C_3, \mathbf{1}\}$, but to include both $\mathbf{1}$ and $\mathbf{0}$ in the picture we consider $\mathbb{C}_+ = \mathbb{C} \cup \{\mathbf{0}, \mathbf{1}\}$.

Given that \mathbb{C}_+ is a lattice, one might ask whether it is a *concept lattice* in the sense defined by FCA. We omit details and limit ourselves to the example of Figure 4. In FCA,

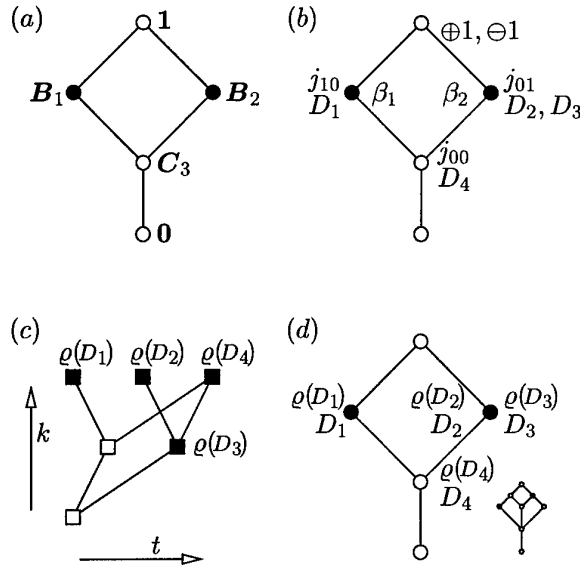


Figure 4: Multiple theories

concepts are certain pairs (*extent, intent*), where a concept's extent is a set of *objects*, and its *intent* is a set of *attributes*. Further, concepts are ordered by a *subconcept/superconcept* relation \leq , and moving up in the lattice of concepts makes the concept's extent (set of objects) bigger, and its intent (set of attributes) smaller. Given our lattice of theories like that of Figure 4, moving up corresponds to expanding the set of theorems: hence, theorems, or descriptions/provable sentences should be treated as objects, in forming extents of theories (concepts). It is appropriate to consider only those descriptions that are axioms of some theories in \mathbb{C} . In the labelled line diagram: see Figure 4(b). We then place descriptions in such a way that they appear at or below the node corresponding to the theory, e.g., in Figure 4 the theory B_2 has the descriptions D_2, D_3 and D_4 , and they can be found by traversing the lattice "down" from the theory's node. Finding concepts' intents is more complicated, but the idea is simple, if you move up from theory to theory, the set of descriptions gets bigger, but "truthfulness" of the theories decreases, and hence we should take as concepts' intents, the sets of "models" of the theories. We first look at Ginsberg's world-based bilattices [Gin88].

Let the set of believed theories \mathbb{B} be seen as a set of *worlds*. We can use the bilattice approach to associate *truth-values* with descriptions. The set of truth-values Γ is given by the power set of \mathbb{B} , $\Gamma = \mathcal{P}(\mathbb{B}) \times \mathcal{P}(\mathbb{B})$ i.e., a truth value is a pair of sets of worlds. A truth-valuation function $\varrho : \mathcal{D} \rightarrow \Gamma$ is given by $\varrho(D) = (U_D, V_D)$, where U_D is a set of worlds where D is *true*, and V_D where it is *false*. In our case it is appropriate to define U_D and V_D as follows. Given $D \in \mathcal{D}$, let $T_D = \text{Cn}(\{D\})$ which is the smallest theory that contains D .

$$U_D = \{B_i \in \mathbb{B} \mid B_i \geq T_D\}$$

$$V_D = \{B_i \in \mathbb{B} \mid B_i \vee T_D = 1\}.$$

Given the example of Figure 4, the truth values for the descriptions are presented in Table 3. Figure 4(c) shows the bilattice-based truth values, together with the bilattice orderings t (truth-ordering on truth-values) and k (information ordering). In Figure 4(d), the truth values are placed next to the corresponding descriptions.

theory	theory's axioms	axioms' truth-value $\varrho(D) = (U_D, V_D)$
B_1	$D_1 = \oplus\{p_1, p_2\}$	$(\{B_1\}, \{B_2\})$
B_2	$D_2 = \ominus\{p_2\}$ $D_3 = \oplus\{p_1, \bar{p}_2\}$	$(\{B_2\}, \{B_1\})$ $(\{B_2\}, \emptyset)$
C_3	$D_4 = \oplus\{p_1\}$	$(\{B_1, B_2\}, \emptyset)$

Table 3: Descriptions and their truth-values

One can note that a single theory can contain several descriptions with different truth values. If this happens, then one can separate such descriptions by presenting a theory as a join of its two subtheories (a small version of such a modified lattice is included in Figure 4(d)). However, what we need is truth-values on theories, rather than single descriptions. Modify the truth valuation function ϱ , so that it takes theories from \mathbb{C} as its arguments, $\varrho : \mathbb{C} \rightarrow \Gamma$, and let it be given by $\varrho(C) = (U_C, V_C)$, where U_C , the set of worlds where C is true, and V_C where the set of worlds C is false are defined to be,

$$U_C = \{B_i \in \mathbb{B} \mid B_i \geq C\}$$

$$V_C = \{B_i \in \mathbb{B} \mid B_i \vee C = 1\}$$

Note that a theory is at most as true and at least as false as its descriptions are. The truth values of the theories of \mathbb{C}_+ are presented in Table 4—trivially, the empty theory $\mathbf{0}$ is true in every believed theory, and false in none, and the reverse applies to $\mathbf{1}$. ($\oplus 1$ and $\ominus 1$ are “axioms” of $\mathbf{1}$).

	theory's truth-value	concept's extent	concept's intent
B_1	$(\{B_1\}, \{B_2\})$	$\{D_1, D_4\}$	$\{j_{10}\}$
B_2	$(\{B_2\}, \{B_1\})$	$\{D_2, D_3, D_4\}$	$\{j_{01}\}$
C_3	$(\{B_1, B_2\}, \emptyset)$	$\{D_4\}$	$\{j_{10}, j_{01}, j_{00}\}$
$\mathbf{0}$	$(\{B_1, B_2\}, \emptyset)$	\emptyset	$\{j_{10}, j_{01}, j_{00}\}$
$\mathbf{1}$	$(\emptyset, \{B_1, B_2\})$	$\{D_1, D_2, D_3, D_4, \oplus 1, \ominus 1\}$	\emptyset

Table 4: Theories as concepts

The problem is that we must *not* be satisfied with truth-values of theories so obtained. Indeed, if we associate a propositional symbol β_i with a statement “the theory (under

consideration) is *true* in B_i ” (i.e., it is below B_i), then we would, for example, associate a sentence $\beta_1 \vee \beta_2$ with C_3 . Propositional models over $\{\beta_i\}_i$ are functions from $\{\beta_i\}_i$ to $\{true, false\}$, so j is a model if $j: \{\beta_i\}_i \rightarrow \{true, false\}$, where we use the convention that, for example, j_{01} denotes a model which satisfies $j(\beta_1) = false$ and $j(\beta_2) = true$. Then, we would get $\{j_{11}, j_{10}, j_{01}\}$ as a set of models associated with C_3 : however no theory can be true in both B_1 and B_2 (because they contradict each other), so the model j_{11} needs to be excluded. In addition, it is not obvious whether j_{00} should be excluded: maybe it is possible that both B_1 and B_2 are “false,” but $C_3 = B_1 \wedge B_2$ nevertheless *is* “true”?

We propose a different method of obtaining *models* for theories. The reason that bilattice-based results are unsatisfactory is that world-based bilattices treat worlds as indistinguishable (see [Gin88], Sections 4 and 7). Our worlds (theories), however are structured, and in the process of deriving models for theories we should make use of this structure.

Let T^\uparrow be a set of total theories, i.e., T^\uparrow is a set of \leq -maximal elements of T . We first propose how to find models for total theories, and then extend this to all theories. Considering our example of Figure 4(a), the theories B_1 , B_2 , C_3 and 0 and their corresponding total theories, i.e., total theories above them, are presented in Figure 5.

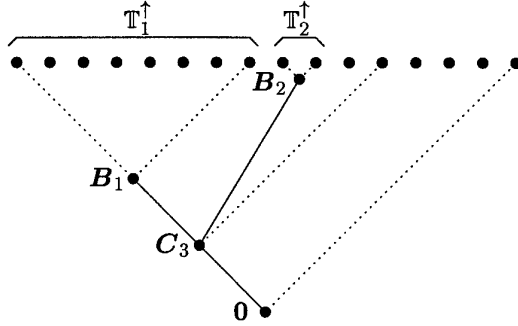


Figure 5: Total theories above B_1 , B_2 , C_3 and 0

Let $T^\uparrow \in T^\uparrow$ be a total theory. Given a believed theory $B_i \in \mathbb{B}$, either T^\uparrow is true in B_i —i.e., $T^\uparrow \geq B_i$ —or T^\uparrow is false in B_i —i.e., $T^\uparrow \vee B_i = 1$; note that no other case is possible, as T^\uparrow is total. Hence, j is a model for T^\uparrow if the truth values j assigns to elements of $\{\beta_i\}$ “agree” with the truthfulness of T^\uparrow in the elements of $\{B_i\}_i$. In other words, j is a model for T^\uparrow if $j(\beta_i) = true$, whenever $T^\uparrow \geq B_i$ (and of course, *false* otherwise). We can now extend this notion of a model for a theory to other theories. We say that j is a model for a theory C iff there is a total theory T^\uparrow above C such that j is a model for T^\uparrow . Figure 5 can be employed to find j -models for B_1 , B_2 , C_3 and 0 . The results are included in Table 4.

Looking at Figure 5, it is clear how to assign a numeric measure to theories. If ω_1 is a *weight* we associate with B_1 , then B_1 would determine a *measure* on total theories, in such a way that the ratio between the measure of a theory in T_1^\uparrow to the measure of a theory in $T^\uparrow \setminus T_1^\uparrow$ is $\omega_1/(1 - \omega_1)$, and the sum of measures on total theories is 1. Similarly,

B_2 would determine “its” measure, and (the two would determine a resulting measure via some quasi-probabilistic method for numerical certainty combination). However, the weights ϖ_i are usually not known. Although some restriction on them can be stipulated (e.g., average weight is greater than $1/2$), or derived: but then, a *maximum entropy principle* can be used to determine the weights. Hence we get a measure on total theories, and this determines a measure on all theories. However, a set of total theories is large, $|\mathbb{T}^\uparrow| = 2^{2^\alpha}$, where $\alpha = |P|$, so such a measure can be difficult to compute: but one can derive a measure on j -models, rather than on total theories.

The above considerations allow us to say that we can derive a preference relation on theories, given just the set \mathbb{B} of believed theories; preference is based on truthfulness and informational value of theories, and thus is a result of how theories agree/reject each other. A method of deriving a numeric measure is sketched which gives a numerical, rather than order-theoretic preference. It is clear that deriving preference relations on objects of interest. In this case on theories, crucial for deciding which information to accept, or, more generally, how to decide *what to prefer*.

4 INTELLIGENCE ANALYSIS AND CONCEPTUAL REASONING

Intelligence analysts are faced with the problem of describing the current, or some predicted future state of the world. To reduce uncertainty, they collect, analyse, combine and fuse information. Information may come from a variety of sources and agencies. The intelligence analyst is normally required to make an assessment, and provide some prediction, when only partial information is held. The various sources and agencies are commonly assessed as being of varying reliabilities, their information of various (believed) accuracies¹.

The information provided may be contradictory, confirmatory, or neither. It may also be either “positive” or “negative” information, ie, “something was observed”, or “something was not observed”. The *conceptual reasoning* framework developed by Nowak [Now97], and described earlier in this report, presents the potential to address these problems with the analysis of intelligence information: to decide what particular subsets (including all) agents agree on, what they disagree on, and on what they reserve judgement. The agents could be simple sensors, Intelligence Agencies with significant processing capability and sources of their own, or anything in between. There may be a mix of types, and there may be a “hierarchy” of sources and agencies, in that some of the agencies may have access to some or all of the information available to some or all other agencies, and may have formed their beliefs after taking these other views and beliefs into account.

The case of a single intelligent agent reasoning with the inputs of a number of simple sensor sources has been dealt with extensively in the literature (on data fusion), and there are a number of possible approaches. The question of whether these approaches can scale to more complex situations remains, we believe, open.

We choose to illustrate the application of Conceptual Reasoning to the domain of Intelligence Analysis with a simple example drawn from the multi-sensor fusion problem of aircraft identification. It is stressed that we are not advocating Conceptual Reasoning as the best approach to this problem; we are using it as a simple example for illustrative purposes.

4.1 Example: Multi-sensor Fusion

Suppose we have three information sources (say, sensors), s_1, s_2, s_3 which can report on the following attributes of observed objects (say, aircraft):

$p_1 = \text{large?}$
 $p_2 = \text{radar?}$
 $p_3 = \text{fast?}$
 $p_4 = \text{high?}$
 ...

¹It is common in Intelligence circles to grade sources in terms of their reliability, from A to F , and information in terms of its (believed) accuracy, from 1 to 6. This is known as the *Admiralty System*.

Our three sensors report as follows:

$$\begin{aligned}
 s_1 \rightarrow D_1 \rightarrow A_1 &= \text{gen}(B_1) = \{\oplus\{p_1, p_2\}\} \\
 s_2 \rightarrow D_2 \rightarrow A_2 &= \text{gen}(B_2) = \{\oplus\{p_1, p_3\}\} \\
 s_3 \rightarrow D_3 \rightarrow A_3 &= \text{gen}(B_3) = \{\ominus\{p_1\}, \oplus\{\bar{p}_1, p_4\}\}
 \end{aligned}$$

Read as “source n produces a description set of observations, from which can be derived a minimal set of logical statements that fully generate the belief that source n holds, and this is ...”.

From these beliefs, we can create a lattice by closure under meet and join (the detailed labelling of the lattice is omitted for clarity).

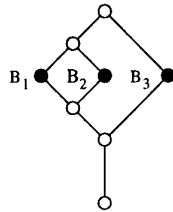


Figure 6: The lattice of beliefs for agents B_1, B_2 and B_3 .

We can see a number of things about the beliefs B_1, B_2 , and B_3 , e.g., there is a possible world (or a number of them) in which the beliefs B_1 and B_2 are compatible (they agree on $\oplus\{p_1\}$), but the beliefs B_1 and B_2 contradict B_3 (who believes $\ominus\{p_1\}$). We can see that B_1 and B_2 agree on something ($\oplus\{p_1\}$), but that $\{B_1$ and $B_2\}$ and B_3 agree only that the world is non-empty.

There are, in principle, a number of states that the world could actually be in². If we had total, accurate, information, we would know which of these states the world was actually in. Our problem is to make some form of statement, or prediction, in light of our partial, perhaps contradictory, information. If we have no information (observations), we can only make the statement that anything is possible (conditioned, perhaps, by some likelihoods derived from past beliefs or observations).

As information is collected, we are able to reduce the size of the set of possible states. We can use the lattice to guide our information collection, as well as to reason with the information that we have.

²It is assumed that states are described at such a level of detail as to make them mutually exclusive, and, of course, exhaustive. If there are n attributes to distinguish objects in the world, there are 2^{2^n} such states, as stated earlier

It is obviously necessary to address contradictions as a matter of priority. We could collect further information, but, once we have a contradiction, we will continue to have it until either:

1. we discount the beliefs of one or more agents entirely, so that contradiction is removed. We might choose this course if we believe one or more agents to be unreliable, or if we believe in the possibility of deliberate deception. In this way, we assign preference to the beliefs of some agents over others, where contradictions occur.
2. we retract some minimal set of beliefs from one or more agents in order to remove contradiction, but leave the remainder of the agents' beliefs intact. This would seem the more appropriate action when we assign no preference over agents, but assign preference over the information itself, i.e., when we hold no intrinsic belief about the reliability ordering of the respective agents, but we do hold some intrinsic belief about the accuracy of the information reported by some agents.

Either way, after the retraction/removal, we reconstruct the lattice and proceed.

Continuing our example, case 1 above results in the lattice at Figure 7, and case 2 above results in Figure 8.

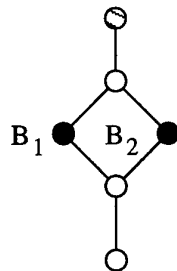


Figure 7: Removing Contradiction by eliminating B_3 influence

We are now dealing with lattices showing no contradiction between agent beliefs. We shall now address the matter of confirmation. In our example, all agents agree the world is non-empty, so it seems reasonable to accept this. B_1 and B_2 agree that there exists an object (or, at least one such object) with attribute p_1 , so we might accept this as confirmation of this “fact” (assuming that s_1 and s_2 are independent sources). Other than these, there is no confirmation, by any agent, of any other agents' beliefs. This touches on the matter of precision versus accuracy, or the granularity of classification. We make a trade off between being more specific, or precise, and being accurate, or correct. In this last example, we can say with confidence that the world is non-empty — all agents agree on this. We can say with less confidence that there is an object with attribute p_1 (two agents agree on this), with still less confidence that there is an object with attribute p_4 ,

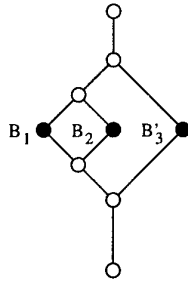


Figure 8: Removing Contradiction by retracting $\ominus\{p_1\}$ from B_3 .

an object with the attributes p_1 and p_2 , and an object with the attributes p_1 and p_3 (only one agent believes each of these).

There are, of course, a number of possible states (at least one) of the world in which all beliefs of all of the agents are true. We have ensured this by removing the contradictions between their beliefs. Using a lattice approach, we are able to examine what theories are possible (or states of the world possible), based on the beliefs of any single agent, any combination of agents, or any shared (i.e., confirmed) beliefs of any combination of agents.

We do this by noting that the lattice we have formed is a sub-lattice of the total lattice of theories in this domain. If we proceed “upwards” in the total lattice, from any point marked by our sub-lattice, we will find, at the penultimate level, all of the possible states, given the beliefs of the starting point.

This is a potentially powerful tool, and it should be noted that this approach generates a strict superset of the union of the possible states generated from each individual agents’ beliefs.

This point is illustrated using the example given earlier in Section 3.4. The believed theories and their axioms appear in table 3. Figure 9 shows the sub-lattice containing the theories. The top row of this figure shows all possible states of the world (K_i). The states consistent with B_1 are ($E_4, E_5, E_6, E_7, E_{12}, E_{13}, E_{14}, E_{15}$). The states consistent with B_2 are (E_8, E_9). The states consistent with both B_1 and B_2 , i.e. (C_3 — where they both confirm each other, is a superset of the union of the states consistent with each theory — it also contains E_{10}) and E_{11}).

4.2 Incorporation of Database Knowledge

Let us suppose that we have a database of knowledge about the possible objects in the world.

We could use this database to “mark” the possible states of the world generated in the above approach, as either “remaining possible”, or “impossible”. Recall there are

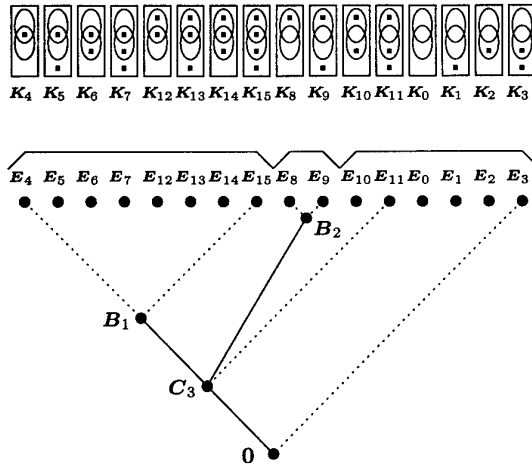


Figure 9: Total theories above B_1 , B_2 , C_3 and 0

Objects	Attributes			
747	+	-	+	...
DC8	+	-	+	...
F16	-	+	+	...
F18	-	+	+	...
Cessna	-	-	-	...
A10	+	+	-	...
B52	+	+	+	...
...

Table 5: Database of Aircraft Attributes

2^{2^n} possible states, where n is the number of attributes. Proceeding with the example illustrated in fig 7 above (where we chose to remove a contradiction by entirely discounting the beliefs of source s_3), we find that there are $2^{2^3} = 256$ possible states! With the database knowledge above, all but 64 of these states are marked “impossible”.

With the aid of suitable notation, the states above each of the beliefs in the lattice can be stated. It is possible to do this by hand for our simple example (see the appendix), but software support would be needed for it to be practical for expected real-world problems. This software has not been developed and could be the subject of a continuation of this research agreement.

Another approach to the incorporation of database knowledge is to treat the database knowledge as another “source”, and to form a lattice of the sources’ beliefs, and the database. The database’s belief, from table 5, would be: $\{\oplus\{p_1, \overline{p_2}, p_3\}, \oplus\{\overline{p_1}, p_2, p_3\}, \oplus\{\overline{p_1}, \overline{p_2}, \overline{p_3}\}$, and so on — remembering, of course, to include \ominus statements for every combination of attributes not found in the database.

The theories “above” any particular combination of beliefs could then be found. Again, this is impractical, for any serious real world problem, without the development of software specific to this purpose, and, again, such software has not been developed for this study.

Despite the limited scope of our study to include only a discussion of the theory of conceptual reasoning and a simple example from the multi-sensor domain, it is natural to make some comments about the implementation of the framework in software. A similar study from the ESM domain [HM86, GS87], encounters a similar multi-agent problem. In that task an Assumption-based Truth Maintenance System (ATMS) [dK86] was used to track the interactions of beliefs. Such a system would no doubt be employed for any implementation of a conceptual reasoning framework. Our claim in this report, however, is that we are setting the scene by developing the appropriate theory that underpins the approach.

5 CONCLUSION

In Section 1 we introduced the scope of the study and discussed how, in general terms, the techniques might be applied to the defence domain. We then presented the preliminary theory of *formal concept analysis* in Section 2.1. Formal concept analysis not only gives the foundations for *conceptual reasoning* but also suggests applications in the defence intelligence domain in its own right.

In Section 3, we presented the theoretical foundations of Conceptual Reasoning. We proposed *partial abstract contexts* as partial models of worlds consisting of objects having attributes. Furthermore, we defined formal systems which allow us to find sentences (descriptions) that follow from a given set, and we related the two, taking into account that the set of contexts \mathbb{K} and the set of theories \mathbb{T} are sets equipped with their information orderings.

We then considered the multiple agent case. Given a set of agents, the set of their *believed theories* is \mathbb{B} , and the resulting lattice is \mathbb{C} , or \mathbb{C}_+ , if we include an empty theory $\mathbf{0}$ and an “inconsistent theory” $\mathbf{1}$. We have shown how one should interpret such structures, and suggested how preference relations on theories can be derived.

As an illustration of the formalism, in Section 4 we showed how conceptual reasoning can be applied in the multi-sensor fusion domain. The complete working of the example is provided as an appendix to this report. We recommend that any further implementation of the framework for real-time fusion and recognition tasks be based on assumption-based truth maintenance or similar technology.

Acknowledgment

The authors acknowledge the assistance of Dr. Martin Oxenham of Wide Area Surveillance Division, DSTO for his careful and thorough reading of an earlier draft of this report. However, any remaining errors are ours alone.

References

- AAA96. Proceedings of the National Conference on Artificial Intelligence. The MIT Press, 1996.
- BFH⁺94. Franz Baader, Enrico Franconi, Bernhard Hollunder, Bernhar Nebel, and Hans-Jürgen Profitlich. An empirical analysis of optimization techniques for terminological representation systems or: Making KRIS get a move on. *Applied Artificial Intelligence. Special Issue on Knowledge Base Management*, 4:109–132, 1994.
- Bur89. Peter Burmeister. Formal concept analysis with a three-valued conceptual logic. TH Darmstadt, 1989.
- Dam95. Marc Damaschek. Gauging similarity with n-grams: Language-independent categorization of text. *Science*, 267:843–848, 10 February 1995.
- DH94. D. Driankov and S. Hellendorn. *Fuzzy Logic and Control*. Springer Verlag, 1994.
- dK86. J. de Kleer. An assumption based tms. *Artificial Intelligence*, 28(1), 1986.
- DP90. B.A. Davey and H.A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, 1990.
- FU95. U. Fayyad and R. Uthurusamy, editors. *The Proceedings of the 1st International Conference on Knowledge and Data Discovery*, Montreal, 1995. AAAI Press.
- Gan96. Bernhard Ganter. Attribute exploration with background knowledge. In *Proceedings, ORDAL96*, 1996.
- GFB⁺92. Michael R. Genesereth, Richard E. Fikes, Daniel Bobrow, Robert MacGregor, Ronald Brachman, John McCarthy, Thomas Gruber, Peter Norvig, Patrick Hayes, Ramesh Patil, Reed Letsinger, Len Shubert, and Vladimir Lifschitz. Knowledge interchange format version 3.0 reference manual. Technical Report Technical Report Logic-92-1, Computer Science Department, Stanford University, Stanford, CA, USA, 1992.
- Gin88. Matthew L. Ginsberg. Multivalued logics: a uniform approach to reasoning in artificial intelligence. *Computational Intelligence*, 4:265–316, 1988.
- Gin93. M. Ginsberg. *Essentials of Artificial Intelligence*. Morgan Kaufmann, 1993.
- GN87. M.R. Genesereth and N.J. Nilsson. *The Logical Foundations of Artificial Intelligence*. Morgan Kaufmann, 1987.
- Gol89. D.E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, 1989.
- GS87. A. Grech and C. Sammut. Qualitative plausible reasoning and assumptions. In *Proceedings of the 3rd Australian Conference on Applications of Expert Systems*, pages 199–218. The NSW Institute of Technology, 1987.

- GW96. Bernhard Ganter and Rudolf Wille. *Formale Begriffsanalyse. Mathematische Grundlagen*. Springer, 1996.
- HM86. S.T. Hood and K.P. Mason. Knowledge-based systems for real-time applications. In *Proceedings of the 2nd Australian Conference on Applications of Expert Systems*, pages 110–137, 1986.
- Hol75. J. Holland. *Adaption in Natural Artificial Systems*. University of Michigan Press, 1975.
- HP91. A. Hertz, J. Krogh and R.G. Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, 1991.
- IJC95. Proceedings of the international conference on artificial intelligence. Morgan Kaufmann, 1995.
- KR96. Proceedings of the international conference on knowledge representation and reasoning. The MIT Press, 1996.
- LJ91. L. Liddy and C.L. Jorgenson. Sublanguage grammer in natural language processing for an expert system. pages 707–717. 1991.
- MA96. P. Martin and L. Alpay. Conceptual structures and structured documents. In *Conceptual Structures: Knowledge Representation as Interlingua*, volume LNAI 1114. Springer Verlag, 1996.
- Mar95. P. Martin. Using the wordnet concept catalog and a relation hierarchy for KA. In *Proceedings of 3rd International Workshop on Peirce*. University of California, 1995.
- Mar96. P. Martin. *Exploitation de graphes conceptuels et de documents structurés et hypertextes pour l'acquisition de connaissances et la recherche d'informations*. PhD thesis, University of Nice - Sophia Antipolis, France, 1996.
- NM96. M.J. Nelson and K.P. Mason. General principles for data fusion systems. In *Proceedings of the First Australian Data Fuzion Symposium*, pages 223–228. IEEE Press, 1996.
- Now97. Chris Nowak. *Conceptual Reasoning*. PhD thesis, Department of Computer Science, 1997.
- Pe93. Sandra Perez and Anthony Sarris (eds). IRDS conceptual schema. Technical Report X3H4/93-196, American National Standards Institute, New York, 1993.
- RN95. S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
- Sow83. J. Sowa. *Conceptual Structures : Information Processing in Mind and Machine*. Addison-Wesley, 1983.

- Vog97. Frank Vogt. Supporting communication in software engineering: an approach based on formal concept analysis, 1997.
- Wil82. R. Wille. Restructuring lattice theory: an approach based on hierarchies of concepts. In *Ordered Sets*, pages 445–470. D. Reidel, Dordrecht, 1982.
- Wil92. R. Wille. Concept lattices and conceptual knowledge systems. *Computers and Mathematical Applications*, 23(6–9):493–515, 1992.
- Wil96. R. Wille. Conceptual structures and multi-contents. In *Proceedings of the 4th International Conference on Conceptual Structures*, volume 987 of *LNAI*, pages 45–56. Springer Verlag, 1996.
- WL90. E. Waltz and J. Llinas. *Multisensor Data Fusion*. Artech House Inc., 1990.
- WMB94. I.H. Witten, A. Moffat, and T. C. Bell. *Managing Gigabytes*. Van Nostrand Reinhold, 1994.

Appendix

The following section demonstrates how the working of the example presented in Section 4 was derived.

Firstly, the total objects over the agencies are computed.

total objects:

object	+valid sentence code
g_000	+{-p1,-p2,-p3} 000
g_002	+{-p1,-p2,p3} 002
g_020	+{-p1,p2,-p3} 020
g_022	+{-p1,p2,p3} 022
g_200	+{p1,-p2,-p3} 200
g_202	+{p1,-p2,p3} 202
g_220	+{p1,p2,-p3} 220
g_222	+{p1,p2,p3} 222

The total theories or total contexts, i.e., subsets of the set of all total objects, can be identified with functions from $\{g_{000}, \dots, g_{222}\}$ to $\{0, 1\}$, i.e. with subsets of the set of all total objects.

{}	00000000	1				*
{g_222}	00000001	0	*	*	*	*
{g_220}	00000010	--> this represents: -----				
{g_220, g_222}	00000011		g	g	g	...
...	...		000	002	020	...
{g_000, ... ,g_222}	11111111					220 222

$B_1 = C_n(\{+{p1,p2}\})$	$gen(B_1) = \{+{p1,p2}\}$
$B_2 = C_n(\{+{p1,p3}\})$	$gen(B_2) = \{+{p1,p3}\}$
$C_3 = B_1 \wedge B_2 =$ $= C_n(\{+{p1}\})$	$gen(C_3) = \{+{p1}\}$
$C_4 = B_1 \vee B_2 =$ $= C_n(\{+{p1,p2},$ $\quad +{p1,p3}\})$	$gen(C_4) = \{+{p1,p2},$ $\quad +{p1,p3}\}$

B B C C (i.e., above B_1, B_2, C_3, C_4?)

is this total theory above the th: 1 2 3 4 ???

{g_222}	00000001	?? ??
{g_220}	002 00000010	?? ??
{g_220, g_222}	00000011	?? ??

.....** above B_1 if 1 there
*. * above B_2 if 1 there
**** above C_3 if 1 there
**** above C_4 if above B_1 and B_2

Below are total theories sorted by their binary codes:

			the sign + or - depending on / whether the given total theory is above B_1 (similarly for B_2, C_3, C_4)
000	00000000	-	- - -
001	00000001	+	+ + +
002	00000010	+	- + -
003	00000011	+	+ + +
004	00000100	-	+ + -
005	00000101	+	+ + +
006	00000110	+	+ + +
007	00000111	+	+ + +
008	00001000	-	- + -
009	00001001	+	+ + +
010	00001010	+	- + -
011	00001011	+	+ + +
012	00001100	-	+ + -
013	00001101	+	+ + +
014	00001110	+	+ + +
015	00001111	+	+ + +
016	00010000	-	- - -
017	00010001	+	+ + +
018	00010010	+	- + -
019	00010011	+	+ + +
020	00010100	-	+ + -
021	00010101	+	+ + +
022	00010110	+	+ + +
023	00010111	+	+ + +
024	00011000	-	- + -
025	00011001	+	+ + +
026	00011010	+	- + -
027	00011011	+	+ + +
028	00011100	-	+ + -
029	00011101	+	+ + +
030	00011110	+	+ + +
031	00011111	+	+ + +
032	00100000	-	- - -
033	00100001	+	+ + +
034	00100010	+	- + -
035	00100011	+	+ + +
036	00100100	-	+ + -
037	00100101	+	+ + +
038	00100110	+	+ + +
039	00100111	+	+ + +
040	00101000	-	- + -
041	00101001	+	+ + +
042	00101010	+	- + -
043	00101011	+	+ + +
044	00101100	-	+ + -
045	00101101	+	+ + +
046	00101110	+	+ + +
047	00101111	+	+ + +
048	00110000	-	- - -
049	00110001	+	+ + +

050	00110010	+ - + -
051	00110011	+ + + +
052	00110100	- + + -
053	00110101	+ + + +
054	00110110	+ + + +
055	00110111	+ + + +
056	00111000	- - + -
057	00111001	+ + + +
058	00111010	+ - + -
059	00111011	+ + + +
060	00111100	- + + -
061	00111101	+ + + +
062	00111110	+ + + +
063	00111111	+ + + +
064	01000000	- - - -
065	01000001	+ + + +
066	01000010	+ - + -
067	01000011	+ + + +
068	01000100	- + + -
069	01000101	+ + + +
070	01000110	+ + + +
071	01000111	+ + + +
072	01001000	- - + -
073	01001001	+ + + +
074	01001010	+ - + -
075	01001011	+ + + +
076	01001100	- + + -
077	01001101	+ + + +
078	01001110	+ + + +
079	01001111	+ + + +
080	01010000	- - - -
081	01010001	+ + + +
082	01010010	+ - + -
083	01010011	+ + + +
084	01010100	- + + -
085	01010101	+ + + +
086	01010110	+ + + +
087	01010111	+ + + +
088	01011000	- - + -
089	01011001	+ + + +
090	01011010	+ - + -
091	01011011	+ + + +
092	01011100	- + + -
093	01011101	+ + + +
094	01011110	+ + + +
095	01011111	+ + + +
096	01100000	- - - -
097	01100001	+ + + +
098	01100010	+ - + -
099	01100011	+ + + +
100	01100100	- + + -
101	01100101	+ + + +
102	01100110	+ + + +
103	01100111	+ + + +
104	01101000	- - + -
105	01101001	+ + + +

106	01101010	+ - + -
107	01101011	+ + + +
108	01101100	- + + -
109	01101101	+ + + +
110	01101110	+ + + +
111	01101111	+ + + +
112	01110000	- - - -
113	01110001	+ + + +
114	01110010	+ - + -
115	01110011	+ + + +
116	01110100	- + + -
117	01110101	+ + + +
118	01110110	+ + + +
119	01110111	+ + + +
120	01111000	- - + -
121	01111001	+ + + +
122	01111010	+ - + -
123	01111011	+ + + +
124	01111100	- + + -
125	01111101	+ + + +
126	01111110	+ + + +
127	01111111	+ + + +
128	10000000	- - - -
129	10000001	+ + + +
130	10000010	+ - + -
131	10000011	+ + + +
132	10000100	- + + -
133	10000101	+ + + +
134	10000110	+ + + +
135	10000111	+ + + +
136	10001000	- - + -
137	10001001	+ + + +
138	10001010	+ - + -
139	10001011	+ + + +
140	10001100	- + + -
141	10001101	+ + + +
142	10001110	+ + + +
143	10001111	+ + + +
144	10010000	- - - -
145	10010001	+ + + +
146	10010010	+ - + -
147	10010011	+ + + +
148	10010100	- + + -
149	10010101	+ + + +
150	10010110	+ + + +
151	10010111	+ + + +
152	10011000	- - + -
153	10011001	+ + + +
154	10011010	+ - + -
155	10011011	+ + + +
156	10011100	- + + -
157	10011101	+ + + +
158	10011110	+ + + +
159	10011111	+ + + +
160	10100000	- - - -
161	10100001	+ + + +

162	10100010	+ - + -
163	10100011	+ + + +
164	10100100	- + + -
165	10100101	+ + + +
166	10100110	+ + + +
167	10100111	+ + + +
168	10101000	- - + -
169	10101001	+ + + +
170	10101010	+ - + -
171	10101011	+ + + +
172	10101100	- + + -
173	10101101	+ + + +
174	10101110	+ + + +
175	10101111	+ + + +
176	10110000	- - - -
177	10110001	+ + + +
178	10110010	+ - + -
179	10110011	+ + + +
180	10110100	- + + -
181	10110101	+ + + +
182	10110110	+ + + +
183	10110111	+ + + +
184	10111000	- - + -
185	10111001	+ + + +
186	10111010	+ - + -
187	10111011	+ + + +
188	10111100	- + + -
189	10111101	+ + + +
190	10111110	+ + + +
191	10111111	+ + + +
192	11000000	- - - -
193	11000001	+ + + +
194	11000010	+ - + -
195	11000011	+ + + +
196	11000100	- + + -
197	11000101	+ + + +
198	11000110	+ + + +
199	11000111	+ + + +
200	11001000	- - + -
201	11001001	+ + + +
202	11001010	+ - + -
203	11001011	+ + + +
204	11001100	- + + -
205	11001101	+ + + +
206	11001110	+ + + +
207	11001111	+ + + +
208	11010000	- - - -
209	11010001	+ + + +
210	11010010	+ - + -
211	11010011	+ + + +
212	11010100	- + + -
213	11010101	+ + + +
214	11010110	+ + + +
215	11010111	+ + + +
216	11011000	- - + -
217	11011001	+ + + +

218	11011010	+ - + -
219	11011011	+ + + +
220	11011100	- + + -
221	11011101	+ + + +
222	11011110	+ + + +
223	11011111	+ + + +
224	11100000	- - - -
225	11100001	+ + + +
226	11100010	+ - + -
227	11100011	+ + + +
228	11100100	- + + -
229	11100101	+ + + +
230	11100110	+ + + +
231	11100111	+ + + +
232	11101000	- - + -
233	11101001	+ + + +
234	11101010	+ - + -
235	11101011	+ + + +
236	11101100	- + + -
237	11101101	+ + + +
238	11101110	+ + + +
239	11101111	+ + + +
240	11110000	- - - -
241	11110001	+ + + +
242	11110010	+ - + -
243	11110011	+ + + +
244	11110100	- + + -
245	11110101	+ + + +
246	11110110	+ + + +
247	11110111	+ + + +
248	11111000	- - + -
249	11111001	+ + + +
250	11111010	+ - + -
251	11111011	+ + + +
252	11111100	- + + -
253	11111101	+ + + +
254	11111110	+ + + +
255	11111111	+ + + +

#####

Now, the same theories sorted by +/-'s, to facilitate how many (and which) theories are above B₁, B₂, C₃, C₄:

001	00000001	+ + + +
003	00000011	+ + + +
005	00000101	+ + + +
006	00000110	+ + + +
007	00000111	+ + + +
009	00001001	+ + + +
011	00001011	+ + + +
013	00001101	+ + + +
014	00001110	+ + + +
015	00001111	+ + + +
017	00010001	+ + + +
019	00010011	+ + + +

021	00010101	+ + + +
022	00010110	+ + + +
023	00010111	+ + + +
025	00011001	+ + + +
027	00011011	+ + + +
029	00011101	+ + + +
030	00011110	+ + + +
031	00011111	+ + + +
033	00100001	+ + + +
035	00100011	+ + + +
037	00100101	+ + + +
038	00100110	+ + + +
039	00100111	+ + + +
041	00101001	+ + + +
043	00101011	+ + + +
045	00101101	+ + + +
046	00101110	+ + + +
047	00101111	+ + + +
049	00110001	+ + + +
051	00110011	+ + + +
053	00110101	+ + + +
054	00110110	+ + + +
055	00110111	+ + + +
057	00111001	+ + + +
059	00111011	+ + + +
061	00111101	+ + + +
062	00111110	+ + + +
063	00111111	+ + + +
065	01000001	+ + + +
067	01000011	+ + + +
069	01000101	+ + + +
070	01000110	+ + + +
071	01000111	+ + + +
073	01001001	+ + + +
075	01001011	+ + + +
077	01001101	+ + + +
078	01001110	+ + + +
079	01001111	+ + + +
081	01010001	+ + + +
083	01010011	+ + + +
085	01010101	+ + + +
086	01010110	+ + + +
087	01010111	+ + + +
089	01011001	+ + + +
091	01011011	+ + + +
093	01011101	+ + + +
094	01011110	+ + + +
095	01011111	+ + + +
097	01100001	+ + + +
099	01100011	+ + + +
101	01100101	+ + + +
102	01100110	+ + + +
103	01100111	+ + + +
105	01101001	+ + + +
107	01101011	+ + + +
109	01101101	+ + + +

110	01101110	+ + + +
111	01101111	+ + + +
113	01110001	+ + + +
115	01110011	+ + + +
117	01110101	+ + + +
118	01110110	+ + + +
119	01110111	+ + + +
121	01111001	+ + + +
123	01111011	+ + + +
125	01111101	+ + + +
126	01111110	+ + + +
127	01111111	+ + + +
129	10000001	+ + + +
131	10000011	+ + + +
133	10000101	+ + + +
134	10000110	+ + + +
135	10000111	+ + + +
137	10001001	+ + + +
139	10001011	+ + + +
141	10001101	+ + + +
142	10001110	+ + + +
143	10001111	+ + + +
145	10010001	+ + + +
147	10010011	+ + + +
149	10010101	+ + + +
150	10010110	+ + + +
151	10010111	+ + + +
153	10011001	+ + + +
155	10011011	+ + + +
157	10011101	+ + + +
158	10011110	+ + + +
159	10011111	+ + + +
161	10100001	+ + + +
163	10100011	+ + + +
165	10100101	+ + + +
166	10100110	+ + + +
167	10100111	+ + + +
169	10101001	+ + + +
171	10101011	+ + + +
173	10101101	+ + + +
174	10101110	+ + + +
175	10101111	+ + + +
177	10110001	+ + + +
179	10110011	+ + + +
181	10110101	+ + + +
182	10110110	+ + + +
183	10110111	+ + + +
185	10111001	+ + + +
187	10111011	+ + + +
189	10111101	+ + + +
190	10111110	+ + + +
191	10111111	+ + + +
193	11000001	+ + + +
195	11000011	+ + + +
197	11000101	+ + + +
198	11000110	+ + + +

199	11000111	+ + + +
201	11001001	+ + + +
203	11001011	+ + + +
205	11001101	+ + + +
206	11001110	+ + + +
207	11001111	+ + + +
209	11010001	+ + + +
211	11010011	+ + + +
213	11010101	+ + + +
214	11010110	+ + + +
215	11010111	+ + + +
217	11011001	+ + + +
219	11011011	+ + + +
221	11011101	+ + + +
222	11011110	+ + + +
223	11011111	+ + + +
225	11100001	+ + + +
227	11100011	+ + + +
229	11100101	+ + + +
230	11100110	+ + + +
231	11100111	+ + + +
233	11101001	+ + + +
235	11101011	+ + + +
237	11101101	+ + + +
238	11101110	+ + + +
239	11101111	+ + + +
241	11110001	+ + + +
243	11110011	+ + + +
245	11110101	+ + + +
246	11110110	+ + + +
247	11110111	+ + + +
249	11111001	+ + + +
251	11111011	+ + + +
253	11111101	+ + + +
254	11111110	+ + + +
255	11111111	+ + + +
002	00000010	+ - + -
010	00001010	+ - + -
018	00010010	+ - + -
026	00011010	+ - + -
034	00100010	+ - + -
042	00101010	+ - + -
050	00110010	+ - + -
058	00111010	+ - + -
066	01000010	+ - + -
074	01001010	+ - + -
082	01010010	+ - + -
090	01011010	+ - + -
098	01100010	+ - + -
106	01101010	+ - + -
114	01110010	+ - + -
122	01111010	+ - + -
130	10000010	+ - + -
138	10001010	+ - + -
146	10010010	+ - + -
154	10011010	+ - + -

162	10100010	+ - + -
170	10101010	+ - + -
178	10110010	+ - + -
186	10111010	+ - + -
194	11000010	+ - + -
202	11001010	+ - + -
210	11010010	+ - + -
218	11011010	+ - + -
226	11100010	+ - + -
234	11101010	+ - + -
242	11110010	+ - + -
250	11111010	+ - + -
004	00000100	- + + -
012	00001100	- + + -
020	00010100	- + + -
028	00011100	- + + -
036	00100100	- + + -
044	00101100	- + + -
052	00110100	- + + -
060	00111100	- + + -
068	01000100	- + + -
076	01001100	- + + -
084	01010100	- + + -
092	01011100	- + + -
100	01100100	- + + -
108	01101100	- + + -
116	01110100	- + + -
124	01111100	- + + -
132	10000100	- + + -
140	10001100	- + + -
148	10010100	- + + -
156	10011100	- + + -
164	10100100	- + + -
172	10101100	- + + -
180	10110100	- + + -
188	10111100	- + + -
196	11000100	- + + -
204	11001100	- + + -
212	11010100	- + + -
220	11011100	- + + -
228	11100100	- + + -
236	11101100	- + + -
244	11110100	- + + -
252	11111100	- + + -
008	00001000	- - + -
024	00011000	- - + -
040	00101000	- - + -
056	00111000	- - + -
072	01001000	- - + -
088	01011000	- - + -
104	01101000	- - + -
120	01111000	- - + -
136	10001000	- - + -
152	10011000	- - + -
168	10101000	- - + -
184	10111000	- - + -

200	11001000	- - + -
216	11011000	- - + -
232	11101000	- - + -
248	11111000	- - + -
000	00000000	- - - -
016	00010000	- - - -
032	00100000	- - - -
048	00110000	- - - -
064	01000000	- - - -
080	01010000	- - - -
096	01100000	- - - -
112	01110000	- - - -
128	10000000	- - - -
144	10010000	- - - -
160	10100000	- - - -
176	10110000	- - - -
192	11000000	- - - -
208	11010000	- - - -
224	11100000	- - - -
240	11110000	- - - -

#####

Now, if we exclude total theories with objects which are disallowed by the database, then we are left with the following total theories:

000	00000000	- - - -
001	00000001	+ + + +
002	00000010	+ - + -
003	00000011	+ + + +
004	00000100	- + + -
005	00000101	+ + + +
006	00000110	+ + + +
007	00000111	+ + + +
008	00001000	- - + -
009	00001001	+ + + +
010	00001010	+ - + -
011	00001011	+ + + +
012	00001100	- + + -
013	00001101	+ + + +
014	00001110	+ + + +
015	00001111	+ + + +
016	00010000	- - - -
017	00010001	+ + + +
018	00010010	+ - + -
019	00010011	+ + + +
020	00010100	- + + -
021	00010101	+ + + +
022	00010110	+ + + +
023	00010111	+ + + +
024	00011000	- - + -
025	00011001	+ + + +
026	00011010	+ - + -
027	00011011	+ + + +
028	00011100	- + + -
029	00011101	+ + + +

030	00011110	+ + + +
031	00011111	+ + + +
128	10000000	- - - -
129	10000001	+ + + +
130	10000010	+ - + -
131	10000011	+ + + +
132	10000100	- + + -
133	10000101	+ + + +
134	10000110	+ + + +
135	10000111	+ + + +
136	10001000	- - + -
137	10001001	+ + + +
138	10001010	+ - + -
139	10001011	+ + + +
140	10001100	- + + -
141	10001101	+ + + +
142	10001110	+ + + +
143	10001111	+ + + +
144	10010000	- - - -
145	10010001	+ + + +
146	10010010	+ - + -
147	10010011	+ + + +
148	10010100	- + + -
149	10010101	+ + + +
150	10010110	+ + + +
151	10010111	+ + + +
152	10011000	- - + -
153	10011001	+ + + +
154	10011010	+ - + -
155	10011011	+ + + +
156	10011100	- + + -
157	10011101	+ + + +
158	10011110	+ + + +
159	10011111	+ + + +

#####

Again, the same theories are sorted by +/-'s, to facilitate the determination of how many (and which) theories are above B₁, B₂, C₃, C₄ (recall that we now consider only those theories which are "allowed" by the database):

001	00000001	+ + + +
003	00000011	+ + + +
005	00000101	+ + + +
006	00000110	+ + + +
007	00000111	+ + + +
009	00001001	+ + + +
011	00001011	+ + + +
013	00001101	+ + + +
014	00001110	+ + + +
015	00001111	+ + + +
017	00010001	+ + + +
019	00010011	+ + + +
021	00010101	+ + + +
022	00010110	+ + + +
023	00010111	+ + + +
025	00011001	+ + + +

027	00011011	+ + + +
029	00011101	+ + + +
030	00011110	+ + + +
031	00011111	+ + + +
129	10000001	+ + + +
131	10000011	+ + + +
133	10000101	+ + + +
134	10000110	+ + + +
135	10000111	+ + + +
137	10001001	+ + + +
139	10001011	+ + + +
141	10001101	+ + + +
142	10001110	+ + + +
143	10001111	+ + + +
145	10010001	+ + + +
147	10010011	+ + + +
149	10010101	+ + + +
150	10010110	+ + + +
151	10010111	+ + + +
153	10011001	+ + + +
155	10011011	+ + + +
157	10011101	+ + + +
158	10011110	+ + + +
159	10011111	+ + + +
002	00000010	+ - + -
010	00001010	+ - + -
018	00010010	+ - + -
026	00011010	+ - + -
130	10000010	+ - + -
138	10001010	+ - + -
146	10010010	+ - + -
154	10011010	+ - + -
004	00000100	- + + -
012	00001100	- + + -
020	00010100	- + + -
028	00011100	- + + -
132	10000100	- + + -
140	10001100	- + + -
148	10010100	- + + -
156	10011100	- + + -
008	00001000	- - + -
024	00011000	- - + -
136	10001000	- - + -
152	10011000	- - + -
000	00000000	- - - -
016	00010000	- - - -
128	10000000	- - - -
144	10010000	- - - -

#####

Conceptual Reasoning and Defence Applications

Peter Deer, Peter W. Eklund and Chris Nowak

(DSTO-RR-0118)

DISTRIBUTION LIST

Number of Copies

AUSTRALIA

DEFENCE ORGANISATION

Task sponsor

Director Defence Intelligence Organisation 1

S&T Program

Chief Defence Scientist)
FAS Science Policy) 1 shared copy
AS Science Corporate Management)
Director General Science Policy Development 1
Counsellor, Defence Science, London Doc Control sheet
Counsellor, Defence Science, Washington Doc Control sheet
Director General Scientific Advisers and Trials) 1 shared copy
Scientific Adviser - Policy and Command)
Navy Scientific Adviser Doc Control sheet
and 1 distribution list
Scientific Adviser - Army Doc Control sheet
and 1 distribution list
Air Force Scientific Adviser 1
Director Trials 1

Aeronautical & Maritime Research Laboratory

Director, 1

Electronics and Surveillance Research Laboratory

Director 1
Chief Information Technology Division 1
Research Leader Command & Control and Intelligence Systems 1
Research Leader Military Computing Systems 1
Research Leader Command, Control and Communications 1
Executive Officer, Information Technology Division Doc Control sheet
Head, Information Architectures Group 1
Head, Information Warfare Studies Group Doc Control sheet
Head, Software Systems Engineering Group Doc Control sheet
Head, Trusted Computer Systems Group Doc Control sheet
Head, Advanced Computer Capabilities Group Doc Control sheet
Head, Computer Systems Architecture Group Doc Control sheet
Head, Systems Simulation and Assessment Group Doc Control sheet
Head, Intelligence Systems Group Doc Control sheet
Head, CCIS Interoperability Lab Doc Control sheet
Head, Command Support Systems Group 1
Head, C3I Operational Analysis Group Doc Control sheet

Head, Human Systems Integration Group	Doc Control sheet
Head, Information Management and Fusion Group	1
Dr M. Nelson	1
Dr M. Oxenham	1
Authors	3
Publications and Publicity Officer, ITD	1
DSTO Library and Archives	
Library Fishermens Bend	1
Library Maribyrnong	1
Library Salisbury	2
Australian Archives	1
Library, MOD, Pyrmont	Doc Control sheet
Capability Development Division	
Director General Maritime Development	Doc Control sheet
Director General Land Development	Doc Control sheet
Director General C3I Development	Doc Control sheet
Intelligence Program	
DGSTA Defence Intelligence Organisation	1
Corporate Support Program (libraries)	
OIC TRS Defence Regional Library, Canberra	1
Officer in Charge, Document Exchange Centre (DEC),	1
US Defence Technical Information Center,	2
UK Defence Research Information Centre,	2
Canada Defence Scientific Information Service,	1
NZ Defence Information Centre,	1
National Library of Australia,	1
Universities and Colleges	
Australian Defence Force Academy	1
Library	1
Head of Aerospace and Mechanical Engineering	1
Deakin University, Serials Section (M list), Deakin University Library,	
Geelong, 3217	1
Senior Librarian, Hargrave Library, Monash University	1
Librarian, Flinders University	1
Other Organisations	
NASA (Canberra)	1
AGPS	1
State Library of South Australia	1
Parliamentary Library, South Australia	1
OUTSIDE AUSTRALIA	
Abstracting and Information Organisations	
INSPEC: Acquisitions Section Institution of Electrical Engineers	1
Library, Chemical Abstracts Reference Service	1
Engineering Societies Library, US	1
Materials Information, Cambridge Scientific Abstracts	1

Documents Librarian, The Center for Research Libraries, US 1

Information Exchange Agreement Partners

Acquisitions Unit, Science Reference and Information Service, UK 1

Library - Exchange Desk, National Institute of Standards and
Technology, US 1

SPARES 10

Total number of copies: 62

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA			1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)		
			N/A		
2. TITLE		3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION)			
Conceptual Reasoning and Defence Applications		Document (U) Title (U) Abstract (U)			
4. AUTHOR(S)		5. CORPORATE AUTHOR			
Peter Deer, Peter W. Eklund and Chris Nowak		Electronics and Surveillance Research Laboratory PO Box 1500 Salisbury SA 5108			
6a. DSTO NUMBER	6b. AR NUMBER	6c. TYPE OF REPORT		7. DOCUMENT DATE	
DSTO-RR-0118	AR-010-397	Research Report		December 1997	
8. FILE NUMBER	9. TASK NUMBER	10. TASK SPONSOR	11. NO. OF PAGES	12. NO. OF REFERENCES	
N8316/5/31	DEF96/167	DDIO	54	36	
13. DOWNGRADING/DELIMITING INSTRUCTIONS			14. RELEASE AUTHORITY		
Not applicable			Chief, Information Technology Division		
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT					
Approved for Public Release					
16. DELIBERATE ANNOUNCEMENT					
No Limitation					
17. CASUAL ANNOUNCEMENT					
Yes					
18. DEFTTEST DESCRIPTORS					
Knowledge Representation, Military Intelligence, Australian Defence Force					
19. ABSTRACT					
This report presents details of a formal technique called <i>conceptual reasoning</i> . Conceptual reasoning is a knowledge representation and reasoning framework for multiple-agent belief revision. Because conceptual reasoning has a particular interest in dealing with <i>contradictory</i> and <i>partial</i> information, it lends itself to applications in Defence Intelligence. We demonstrate the practical application of conceptual reasoning in a particular defence intelligence domain: multi-sensor fusion.					