NORTH ATLANTIC TREATY ORGANIZATION

SCIENCE AND TECHNOLOGY ORGANIZATION

AC/323(IST-111)TP/1022

www.sto.nato.int

STO TECHNICAL REPORT **TR-IST-ET-111**

Knowledge Representation and Reasoning – A Review of the State of the Art and Future Opportunities

(Représentation des connaissances et raisonnement – revue de l'état de la technique et opportunités futures)

A report from NATO IST ET-111.

Published June 2022

Distribution and Availability on Back Cover

NORTH ATLANTIC TREATY ORGANIZATION

SCIENCE AND TECHNOLOGY ORGANIZATION

AC/323(IST-111)TP/1022

www.sto.nato.int

STO TECHNICAL REPORT **TR-IST-ET-111**

Knowledge Representation and Reasoning – A Review of the State of the Art and Future Opportunities

(Représentation des connaissances et raisonnement – revue de l'état de la technique et opportunités futures)

A report from NATO IST ET-111.

The NATO Science and Technology Organization

Science & Technology (S&T) in the NATO context is defined as the selective and rigorous generation and application of state-of-the-art, validated knowledge for defence and security purposes. S&T activities embrace scientific research, technology development, transition, application and field-testing, experimentation and a range of related scientific activities that include systems engineering, operational research and analysis, synthesis, integration and validation of knowledge derived through the scientific method.

In NATO, S&T is addressed using different business models, namely a collaborative business model where NATO provides a forum where NATO Nations and partner Nations elect to use their national resources to define, conduct and promote cooperative research and information exchange, and secondly an in-house delivery business model where S&T activities are conducted in a NATO dedicated executive body, having its own personnel, capabilities and infrastructure.

The mission of the NATO Science & Technology Organization (STO) is to help position the Nations' and NATO's S&T investments as a strategic enabler of the knowledge and technology advantage for the defence and security posture of NATO Nations and partner Nations, by conducting and promoting S&T activities that augment and leverage the capabilities and programmes of the Alliance, of the NATO Nations and the partner Nations, in support of NATO's objectives, and contributing to NATO's ability to enable and influence security and defence related capability development and threat mitigation in NATO Nations and partner Nations, in accordance with NATO policies.

The total spectrum of this collaborative effort is addressed by six Technical Panels who manage a wide range of scientific research activities, a Group specialising in modelling and simulation, plus a Committee dedicated to supporting the information management needs of the organization.

- AVT Applied Vehicle Technology Panel
- HFM Human Factors and Medicine Panel
- IST Information Systems Technology Panel
- NMSG NATO Modelling and Simulation Group
- SAS System Analysis and Studies Panel
- SCI Systems Concepts and Integration Panel
- SET Sensors and Electronics Technology Panel

These Panels and Group are the power-house of the collaborative model and are made up of national representatives as well as recognised world-class scientists, engineers and information specialists. In addition to providing critical technical oversight, they also provide a communication link to military users and other NATO bodies.

The scientific and technological work is carried out by Technical Teams, created under one or more of these eight bodies, for specific research activities which have a defined duration. These research activities can take a variety of forms, including Task Groups, Workshops, Symposia, Specialists' Meetings, Lecture Series and Technical Courses.

The content of this publication has been reproduced directly from material supplied by STO or the authors.

Published June 2022

ISBN 978-92-837-2342-4

Table of Contents

List of Figures

Glossary

Ambiguity	The quality of being open to more than one interpretation; inexactness [1].
Artificial Intelligence	Intelligence exhibited by machines. A branch of computer science dealing with the simulation of intelligent behaviour in computers. The capability of a machine to imitate intelligent human behaviour. The ability of machines to match humans in terms of learning, reasoning, planning and acting in complex cyber-physical environments. The study of "intelligent agents": any device that perceives its environment and takes actions that maximize its chance of success at some goal. When a machine mimics "cognitive" functions that humans associate with other human minds, such as "learning" and "problem solving". The science of making computers do things that require intelligence when done by humans. The ability of a computer program or a machine to think and learn. A field of study which tries to make computers "smart".
Backward Chaining	Starting with something one wants to prove, finding implication sentences that would allow him/her to conclude it, and then attempting to establish their premises in turn [2].
Completeness	An inference procedure is complete if it can find a proof for any sentence that is entailed [2].
Data	Individual observations, measurements, and primitive messages from the lowest level of abstraction. Human communication, text messages, electronic queries, or scientific instruments that sense phenomena are the major sources of data. The term evidence (data that is determined to be relevant) is frequently used to refer to elements of data [3].
Deduction	Reasoning about premises to derive conclusions [3].
Domain	A domain is a body of knowledge [4]. In knowledge representation, a domain is a section of the world [2].
Expert	A person who has (or is recognized by peers as having) expertise in a certain area [5]. In general, the term expert connotes both specialization in narrow problem-solving areas or tasks and substantial competence [4], [6]. An expert can solve problems that most people cannot solve, or can solve them more efficiently (but not as cheaply) [5].

Expertise	Expertise is a specialized type of knowledge that is known only to a few. It is not commonly found in public sources such as books and papers. Instead, expertise is the extensive, task-specific and implicit knowledge of the expert that is acquired from training, reading, and experience [5].
Expert System	An intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution [5]. The user supplies facts or other information to the expert system and receives expert advice or expertise in response [5].
Forward Chaining	Reasoning from facts to the conclusion(s) resulting from those facts [5].
Graph	Mathematical structure used to model pairwise relations between objects [7]. A graph in this context is made up of vertices (also called nodes or points) which are connected by edges (also called links or lines).
Inference Engine	A software component that reasons its way to the solutions of problems, with its search guided by the contents of a knowledge base. It must include provisions for setting goals, representing and recording intermediate results, and managing memory and computational resources [4].
Inference Procedure	Given a knowledge base KB, an inference procedure can generate new sentences that purport to be entailed by KB, or, given a knowledge base KB and another sentence, an inference procedure can report whether or not the sentence is entailed by the KB [2].
Information	Organized sets of data. The organization process may include sorting, classifying, or indexing and linking data to place data elements in relational context for subsequent searching and analysis [3].
Information Fusion	The process of utilizing one or more information sources over time to assemble a representation of aspects of interest in an environment [8].
Interpretation	Say what fact a sentence corresponds to. A systematic relationship between sentences and facts [2]. A sentence is true under a particular interpretation if the state of affairs it represents is the case [2].
Knowledge	A relation between a knower and a proposition [9]. The fact or condition of knowing something with a considerable degree of familiarity through experience, association or contact [10]. A dynamic human process of justifying human belief toward the truth [11]. The codified experience of agents. Codified emphasizes that knowledge is written. Experience emphasizes that knowledge is created and used in experiential situations. Agents undergo experiences [4].

Information once analysed, understood, and explained is knowledge, or foreknowledge (predictions or forecasts). Understanding information provides:

1) A degree of comprehension of both the static and dynamic relationships of the objects of data;
2) The ability to model structures; and
3) Past (and future) behaviour of those objects.

Knowledge includes both static content and dynamic processes [3].

Knowledge Acquisition	The process of collecting, extracting, transferring, accumulating, structuring, transforming and organizing knowledge (e.g., problem-solving expertise) from one or more knowledge sources (human experts, books, documents, sensors, or computer files) for constructing or expanding a knowledge base [6].
Knowledge Base	The organized repository for the collection of knowledge related to a domain and used for understanding, formulating, and solving problems in a knowledge-based system [4]. A collection of symbolic structures representing what a system believes and reasons with during its operation [9]. A set of representation of facts about the world [2]. Contains the knowledge with which the inference engine draws conclusions [5].
Knowledge-Based System	A computer system that represents and uses knowledge to carry out a task [4]. A knowledge-based system is composed of a knowledge base and an inference mechanism. It operates by storing sentences about the world in its knowledge base, using the inference mechanism to infer new sentences, and using them to decide what action to take [2].
Knowledge Engineer	Someone who investigates a particular domain, determines what concepts are important in that domain, and creates a formal representation of the objects and relations in the domain [2].
Knowledge Engineering	The process of building a knowledge base. It deals with knowledge acquisition, knowledge representation, knowledge validation, inferencing, explanation and justification, and maintenance [2].
Knowledge Graph	A programmatic way to model a knowledge domain with the help of subject-matter experts, data interlinking, and machine learning algorithms [12]. A knowledge graph acquires and integrates information into an' ontology and applies a reasoner to derive new knowledge [13].
Knowledge Representation	The field of study concerned with using formal symbol to represent a collection of propositions believed by some putative agent [9]. Expressing knowledge in computer-tractable form, such that it can be exploited [2]. Finding a language in which to encode knowledge so that the machine can use it [14].

Knowledge Representation and Reasoning	The area of artificial intelligence concerned with how knowledge can be represented symbolically and manipulated in an automated way by reasoning programs [9].
Logic	The study of entailment relations [9]. A language, Provided the syntax and semantics are defined precisely [2].
Meaning	What a sentence states about the world, that the world is this way and not that way [2].
Model	A world in which a sentence is true under a particular interpretation is called a model of that sentence under that interpretation [2].
Monotonicity	A logic is monotonic if when we add some new sentences to the knowledge base, all the sentences entailed by the original KB are still entailed by the new larger knowledge base [2].
Ontology	A formal, explicit specification of a shared conceptualization [15].
Ontological Engineering	The set of activities that concern the ontology development process, the ontology life cycle, the methods and methodologies for building ontologies, and the tool suites and languages that support them [16].
Proposition	The idea expressed by a simple declarative sentence [9]. An abstract entity that can be true or false, right or wrong, factual or nonfactual [9]. A classification of all the different ways one can imagine the world to be [9].
Reasoning / Inference	Using the faculty of reason so as to arrive at conclusions [10]. To discover, formulate, or conclude by the use of reason [10]. To derive as a conclusion from facts or premises [10]. The formal manipulation of symbols representing a collection of believed propositions to produce representations of new ones [9]. Determining what follows from what the knowledge base has been told [2]. Any process by which conclusions are reached [2].
Representation	A relationship between two domains [9]. The assignment of meaning to symbols [4].
Semantics	Specify what the well-formed expressions are supposed to mean [9]. Determines the facts in the world to which the sentences refer. Each sentence makes a claim about the world [2].
Situation Analysis	A process, the examination of a situation, its elements, and their relations, to provide and maintain a product, i.e., a state of situation awareness for the decision maker(s) [17].

Situation Awareness	The perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future [18].
Symbol	A character or group of characters taken from some predetermined alphabet [9].
Syntax	Specify which groups of symbols, arranged in what way, are to be considered properly formed [9]. Describes the possible configurations that can constitute sentences [2].
Taxonomy	The branch of science concerned with classification, especially of organisms; systematics [19].
Uncertainty	Uncertainty refers to epistemic situations involving imperfect or unknown information [20]. It applies to predictions of future events, to physical measurements that are already made, or to the unknown. Uncertainty arises in partially observable and/or stochastic environments, as well as due to ignorance, indolence, or both.
Validity (Tautology)	A sentence is valid or necessarily true if and only if it is true under all possible interpretations in all possible worlds, i.e., regardless of what it is supposed to mean and regardless of the state of affairs in the universe being described [2].

REFERENCES

[1] https://www.google.com/search?q=define+ambiguity&rlz=1C1CAFA_enCA859CA859&oq=define+ambiguity&aqs=chrome..69i57j0l7.4479j1j8&sourceid=chrome&ie=UTF-8; Accessed May 2020.

[2] Russel, S. and Norvig, P.(1995). Artificial Intelligence, A Modern Approach. Prentice Hall, New Jersey.

[3] Waltz, E. (2003). Knowledge Management in the Intelligence Enterprise. Artech House, Norwood, MA.

[4] Stefik, M. (2014). Introduction to Knowledge Systems. Morgan Kaufmann.

[5] Giarratano, J. (2005). Expert Systems: Principles and Programming. Thomson, Switzerland.

[6] Turban, E. and Aronson, J.E. (1998). Decision Support Systems and Intelligent Systems. New Jersey, Prentice Hall.

[7] Wikipedia. Graph Theory. https://en.wikipedia.org/wiki/Graph_theory#Graph; Accessed May 2020.

[8] Lambert, D.A., Grand Challenges of Information Fusion, Proceedings of the Sixth International Conference on Information Fusion, Cairns, Australia, 2003, pp. 213-219.

[9] Brachman, R. and Levesque, H. (2004). Knowledge Representation and Reasoning. Morgan Kaufmann.

[10] Webster's (1981). Webster's Third New International Dictionary. Springfield MA, Merriam-Webster.

[11] Nonaka, I. (2007). The Knowledge-Creating Company. Harvard Business Review.

[12] Aijal, J. (2019). What Is a Knowledge Graph and How Does One Work? Retrieved May, 2020. https://thenextweb.com/news/what-is-a-knowledge-graph-and-how-does-one-work.

[13] Ehrlinger, L. and Wöß, W. (2016). Towards a Definition of Knowledge Graphs. 1st International Workshop on Semantic Change & Evolving Semantics (SuCCESS16). Leipzig, Germany, 1695.

[14] Ginsberg, M. (1993). Essentials of Artificial Intelligence, Morgan Kaufmann.

[15] Studer, R., Benjamins, V.R. and Fensel, D. (1998). Knowledge Engineering: Principles and Methods. IEEE Transactions on Data and Knowledge Engineering 25, pp. 161-197.

[16] Gómez-Pérez, A., Fernández-López, M. and Corcho, O. (2004). Ontological Engineering, Springer.

[17] Roy, J. (2001). From Data Fusion to Situation Analysis. the Fourth International Conference on Information Fusion (FUSION 2001). Montreal, Canada.

[18] Endsley, M.R. (1995). Toward a Theory of Situation Awareness in Dynamic Systems. Human Factors(1), p. 32.

[19] https://www.google.com/search?q=define+taxonomy&rlz=1C1CAFA_enCA859CA859&oq=define+taxonomy&aqs=chrome..69i57j0l7.5118j1j8&sourceid=chrome&ie=UTF-8; Accessed May 2020.

[20] Wikipedia. Uncertainty. https://en.wikipedia.org/wiki/Uncertainty; Accessed May 2020.

IST-ET-111 Membership List

CHAIR

Dr. David BARBER
Dstl
UNITED KINGDOM
Email: dbarber@dstl.gov.uk

MEMBERS

Dr. Claire BONIAL
U.S. Army Research Laboratory (ARL)
UNITED STATES
Email: Claire.N.Bonial.civ@mail.mil

Dr. Paul CRIPPS
Dstl
UNITED KINGDOM
Email: pcripps@dstl.gov.uk

Dr. Forrest HARE
Defense Intelligence Agency
UNITED STATES
Email: Forrest.b.hare@saic.com

Dr. Marielle MOKHTARI
Defence R&D Valcartier
CANADA
Email: marielle.mokhtari@forces.gc.ca

Dr. Chris MOWAT
Dstl
UNITED KINGDOM
Email: cfmowat@dstl.gov.uk

Mr. Jean ROY
Defence R&D Valcartier
CANADA
Email: jean.roy@drdc-rddc.gc.ca

Dr. Hans-Christian SCHMITZ
Fraunhofer FKIE
GERMANY
Email: hans-christian.schmitz@fkie.fraunhofer.de

Mr. John SWEET
National Geospatial-Intelligence Agency (NGA) (Contractor)
UNITED STATES
Email: John.W.Sweet.ctr@nga.mil

PANEL/GROUP MENTOR

Prof. Bhopinder MADAHAR
Dstl
UNITED KINGDOM
Email: bkmadahar@dstl.gov.uk

Knowledge Representation and Reasoning – A Review of the State of the Art and Future Opportunities

(STO-TR-IST-ET-111)

Executive Summary

This report presents the findings of NATO Information Systems Technology (IST) Exploratory Team 111 (ET-111). ET-111 was formed to share understanding across NATO nations on the status of Knowledge Representation and Reasoning (KRR) in order to understand the state of the art and to consider if future collaborative activities would be beneficial.

In support of high level data fusion there is now a foundational need for information and knowledge to be understandable by humans and machines. Knowledge representation is the expression of knowledge in computer-tractable form in order for it to be exploited. A key but not sole reason for doing this is so that the knowledge can be reasoned over. Knowledge-based systems might also be referred to as symbolic AI and rule-based AI, and have been an active area of research for over five decades. As such, it might be considered by some as ‘old-school’ AI, differing from the algorithmic, machine-learning-based approaches to AI that have grown in prominence in recent years (and which are known to suffer from problems with explainability and generalisation). In the age of ‘big data’, knowledge representation and reasoning provides an avenue for the exploitation of data which is flexible, explainable, and grounded in human knowledge.

The first aim of this review was to provide a technical introduction to the field of knowledge representation and reasoning. Providing the reader with knowledge of key concepts – to nurture understanding – will enable an appreciation of the capabilities of knowledge systems. The second aim is to provide – by example – a working grasp of the processes one must utilise in order to create a knowledge system, and how such systems can be used in a military context to solve real-world problems. An understanding of such real-world problems to which knowledge systems are best applied should facilitate successful implementation and integration of KRR with NATO systems and doctrine.

In this report, we begin by discussing some of the challenges for NATO member nations, and how knowledge representation and reasoning in NATO may be expected to have an influence on these areas. We then provide a summary of the technical aspects of knowledge representation, knowledge engineering and methods for reasoning. We discuss specific examples of knowledge representation such as the MIP Information Model (MIM), the Rich Event Ontology (REO), OPIS, and the Defense Intelligence Core Ontology (DICO). We also describe the WISDOM R&D platform and the Intelligent Situational Awareness (INSANE) framework as examples of using knowledge representation to support sense making.

Following this we review wider research, including how text analytics can support the extraction of knowledge from reports and other sources of text, work on causality and the problems of explainability and trust in reasoning systems.

Finally, we conclude with a summary of the findings of the report and the implications this for the Alliance presenting key recommendations for further work:

- Recommendation 1 – The NATO STO sponsors a technical activity to demonstrate the complementary use of symbolic and sub-symbolic methods and their benefit to improved decision making.
- Recommendation 2 – The NATO STO sponsors a virtual lecture series/workshop to increase the awareness of KRR technologies in the science and operational sectors of NATO, in order to provide a catalysis for further skills development in this area.
- Recommendation 3 – The NATO STO sponsors a dedicated Exploratory Team to consider specific interests in causal modelling and its application to knowledge-based systems, as a possible precursor to future practical demonstrations under activities such as that against Recommendation 1.

Représentation des connaissances et raisonnement – revue de l'état de la technique et opportunités futures

(STO-TR-IST-ET-111)

Synthèse

Le présent rapport expose les conclusions de l'équipe exploratoire 111 (ET-111) de la commission OTAN sur la technologie des systèmes d'information (IST). L'ET-111 a été constituée pour confronter les différents points de vue des pays de l'OTAN sur le statut de la représentation des connaissances et du raisonnement (KRR), afin de comprendre l'état de la technique et d'étudier si de futures activités en collaboration seraient bénéfiques.

La fusion de données à haut niveau s'inscrit désormais dans un contexte où il est fondamental que les humains et les machines puissent comprendre les informations et les connaissances. La représentation des connaissances est l'expression des connaissances sous une forme pouvant être exploitée par un ordinateur. La possibilité de raisonner à partir de ces connaissances est l'une des raisons essentielles de cette représentation, mais non la seule. Les systèmes basés sur les connaissances peuvent également être appelés IA symbolique et IA basée sur des règles et constituent un domaine de recherche active depuis plus de cinquante ans. En tant que tels, ils pourraient être considérés comme une IA « à l'ancienne », se distinguant des approches algorithmiques de l'IA basées sur l'apprentissage automatique, qui ont pris de l'importance ces dernières années. À l'ère des « données massives », la représentation des connaissances et le raisonnement constituent un moyen d'exploiter les données qui est flexible, explicable et enraciné dans les connaissances humaines.

Le premier objectif de la présente revue est de fournir une introduction technique au domaine de la représentation des connaissances et du raisonnement. La connaissance des concepts essentiels, favorisant la compréhension, permettra au lecteur d'apprécier les capacités des systèmes de connaissances. Le deuxième objectif est de fournir, au moyen d'exemples, une appréhension pratique i) des processus à utiliser pour créer un système de connaissances et ii) de la manière dont ces systèmes peuvent servir dans un contexte militaire pour résoudre des problèmes réels. La compréhension des problèmes réels auxquels les systèmes de connaissances sont le plus efficacement appliqués devrait faciliter la mise en œuvre et l'intégration de la KRR dans les systèmes et la doctrine de l'OTAN.

Dans ce rapport, nous commençons par discuter de quelques défis qui se présentent aux pays de l'OTAN et de la façon dont la représentation des connaissances et le raisonnement au sein de l'OTAN pourraient avoir une influence sur ces questions. Nous résumons ensuite les aspects techniques de la représentation des connaissances, du génie de la connaissance et des méthodes de raisonnement. Nous discutons d'exemples particuliers de représentation des connaissances, tels que le modèle d'information du MIP (MIM), l'ontologie d'événements riche (Rich Event Ontology, REO), OPIS et l'ontologie essentielle du renseignement de la défense (Defense Intelligence Core Ontology, DICO). Nous décrivons également la plateforme de R&D WISDOM et le cadre de connaissance intelligente de la situation (Intelligent Situational Awareness, INSANE) en tant qu'exemples d'utilisation de la représentation des connaissances pour faciliter l'attribution d'un sens.

À la suite de cela, nous passons en revue la recherche dans son ensemble, notamment la manière dont l'analyse de texte peut favoriser l'extraction de connaissances à partir de rapports et d'autres sources de texte, le travail sur la causalité et les problèmes d'explicabilité et de confiance dans les systèmes de raisonnement.

Enfin, nous résumons les conclusions du rapport et leurs implications pour l'Alliance, en présentant des recommandations clés pour les travaux ultérieurs :

- Recommandation n° 1 – Que la STO de l'OTAN parraine une activité technique pour démontrer l'utilisation complémentaire des méthodes symboliques et sous-symboliques et leur avantage sur le plan de l'amélioration des décisions.
- Recommandation n° 2 – Que la STO de l'OTAN parraine une série de conférences ou un séminaire virtuels, afin d'augmenter la connaissance des technologies de KRR dans les secteurs scientifique et opérationnel de l'OTAN, de manière à catalyser le développement ultérieur de compétences dans ce domaine.
- Recommandation n° 3 – Que la STO de l'OTAN parraine une équipe exploratoire étudiant l'intérêt de la modélisation causale et son application aux systèmes basés sur les connaissances, comme préalable à d'éventuelles démonstrations pratiques dans le cadre d'activités telles que celles de la recommandation n° 1.

Chapter 1 – INTRODUCTION

1.1 INFORMATION 'IN WAR'

"…a wealth of information creates a poverty of attention…"

Herbert A. Simon, 1971, economist, psychologist and Nobel Prize winner [1].

With an ever increasing number of capabilities, sensors, feeds and other data one of the most pressing challenges facing defence is the ability to reliably and quickly sift, fuse and act on the most pertinent observations and information. The importance of information is manifest in all of NATO's strategic priorities [2]. Russia's threat to Euro-Atlantic security is one based on disinformation intended to undermine strategic relationships (e.g., European Union, NATO, etc.); the fight against terrorism, in all its forms and manifestations, is now predicated on being able to connect both classified and open source material to identify connections and behaviours against which action can be taken, and; the cyber threat is one principally fought in information space.

NATO's acquisition of the Alliance Ground Surveillance (AGS) system represents a significant enhancement to NATO's capability to provide rich data feeds in support of its future operations [3]. But it is recognised, in concepts such as the UK's "Information Advantage", that a real advantage can now only be achieved by the timely and effective fusion of such data feeds.

Of course, the challenge of dealing with information overload is not limited to defence. The worlds of finance, advertising and engineering, to name a few, are embracing the opportunities to improve decision making, target services and increase the pace at which new solutions can be delivered. In recent years the potential of applying Machine Learning (ML) approaches to these challenges has caught the imagination of the public, investors and senior leaders around the world. As a result such ML methods are now demonstrating their potential against defence challenges, including object detection and labelling in imagery and video feeds, text analysis to extract entities and relationships and speech detection and translation. Building on the revolution in computational power, data availability and access to computational frameworks the explosion of interest in data and Artificial Intelligence (AI) is now considered a turning point for how Defence works, and is recognised as being comparable to the first aircraft or nuclear weapons [4].

As a result, Defence capabilities will increasingly be able to deal with the most critical information streams, saving analyst time and increasing their capacity to quickly develop and retain situational awareness. Nevertheless, as the attention of human analysts also becomes stretched by the activities of operating and warfighting in a time of persistent competition [5], there remains a need to continually improve their ability to connect subtle but significant observations across multiple domains. For example, as operations look to routinely apply full spectrum effects, the interconnection of observations between the physical, social and cyber domains will be increasingly important, but such connections may not be easily discernible without supporting analytical capabilities. Importantly it will be vital to connect such observations to past knowledge and the inherent expertise of those involved and the experience of those before them.

1.2 UNDERSTANDING AND INFORMATION FUSION

Complexity has always existed in natural and biological realms. However, with advances in sciences and technology, humanity is now capable of building artefacts whose complexity approaches those of life itself. There is a need to use advanced methods to tackle this complexity.

This complexity stems from tremendous increases in speed, density, and spatial scope of data, alongside a coupling between an ever increasing range of elements, some natural, and many synthetic. Approaches to cope with the complexity of situations is a core defence challenge. We are faced with the rapid evolution of technology providing more data, information and capabilities combined with challenging terrains e.g., urban environments and the "human terrain" involving insurgents, mixed populations, non-government organisations and failed states. These complex situations actually, more than ever, require timely decisions to overmatch the threat, and decision quality will always be closely associated to the level of situation comprehension. That comprehension, challenged by such rapid operational and technological changes, requires new approaches for the better and faster disentangling of complex situations. While the Human mindset is still deeply rooted in the classical concept of reductionism, in which a problem is solved by decomposing it into sub-problems, it is now recognised that approaching complexity in an effective way cannot be isolated from a reductionism approach.

National concepts such as the UK's Information Advantage [6] concept have sought to catalyse the role of information in defence operations, emphasising the need to innovate or risk "withering" and losing pace with adversaries. The US Augmenting Intelligence with Machines (AIM) initiative [7] also offers a strategic viewpoint and emphasises the role of AI and ML in future intelligence capabilities. Of particular note to IST-ET-111 the AIM initiative makes a priority of basic research advances in representing knowledge.

The high level technical challenge in achieving both low and high level data fusion has been well defined over the years by the JDL Fusion Model [8]. However, comprehensive solutions, in particular for high level data fusion, remain absent and the subject of ongoing research and development.

The challenge of information fusion extends to almost all aspects of defence from logistics to personnel management, platform maintenance and medical treatment. However, in a NATO context with a focus on coalition operations such as those in Afghanistan it is perhaps most pertinent to consider the challenge of achieving situational awareness to support Command and Control (C2) and intelligence functions. Such activities are characterised by the need to:

- Draw together primary observations and less tangible information and knowledge (hard/soft fusion);
- Bring together information from multiple domains and often at multiple classifications;
- Make decisions on limited information and at high tempo; and
- Deal with uncertainty, ambiguity and ever changing information.

The primary authors and consumers of the outputs of C2 and intelligence activities are human analysts and warfighters, but as NATO forces make more use of automation and autonomous systems the role of machines in supporting, augmenting and exploiting underpinning situational awareness and high level fusion activities will be increasingly important to consider.

1.3 THE ROLE OF KNOWLEDGE REPRESENTATION AND REASONING

In support of such high level fusion there is now a foundational need for information and knowledge to be understandable by humans and machines. By doing this it becomes possible to apply machine reasoning (inferencing) methods, which apply rules and formal logic to available data in order to offer higher order deductions. Knowledge representation is the expression of knowledge in computer-tractable form in order for it to be exploited. A key but not sole reason for doing this is so that the knowledge can be reasoned over. Knowledge-based systems are also referred to by terms such as symbolic AI and rule-based AI, and have been an active area of research for over five decades.

Knowledge representation and the desire to reason against this knowledge lies at the heart of three of NATO's seven Emerging and Disruptive Technologies (EDTs): data, AI and autonomy, each of which overlap and support the other [9].

Of course, it is the most recent interest in "machine learning" (also known as sub-symbolic) methods, focusing on computational approaches such as neural networks, which has been at the forefront of the public narrative of AI, often seeing the terms ML and AI used interchangeably. Such systems have proven their value in multiple applications such as product recommender systems, the prediction of traffic patterns and loan approvals. The use of ML to support content-based analytics of multiple data types (imagery, video, text and social media) was the core interest of the now completed IST-RTG-144 (Multi-content analytics). The team clearly demonstrated the potential for individual modalities to be analysed, and the opportunity to combine those analytics within the wider intelligence cycle, but it did not consider the automated/semi-automated fusion or reasoning against observations to support the goals of high level fusion.

ML approaches generally capitalise on large volumes of data to develop models to associate outputs to inputs. For some classes of task, such as imagery labelling, ML is a proven application, but even leaders in AI systems are cautious, if not critical in its success [10]. The representation and inclusion of knowledge with ML approaches (so called neuro-symbolic approaches) could be a potential step in increasing the robustness and performance of future solutions.

As such, a fundamental step towards this goal is to establish effective knowledge representation (symbolic representation) that can be used by future hybrid systems. Symbolic methods may be more adept at dealing with sparse data, support enhanced explainability and incorporate past human knowledge, and using computational methods which excel at pattern recognition and data clustering/classification problems. However, if such approaches/technologies are to support future coalition operations, joined up effort is required. This includes the:

- Development of domain specific ontologies (defined vocabularies for specific domains);
- Deployment and assessment of inferencing capabilities;
- Building effective architectures for event driven processing;
- Handling uncertainty and ambiguity in observations;
- Information sharing and observation provenance; and
- Approaches for federated deployment and coping with scale.

1.4 AIMS AND OBJECTIVES OF IST-ET-111

It is in this context that NATO IST-ET-111 "Knowledge Representation and Reasoning" was proposed, with the aim of pooling understanding across interested NATO partners on the status of KRR in order to understand the state of the art and to consider if future activities might be required.

By establishing the current state of the art and the technical capabilities across NATO nations the IST-ET-111 team hopes this report will support a conversation around the most effective ways to achieve effective human-machine teams. The establishment of an Exploratory Team is also a step in identifying how other NATO nation science and technology activities might be leveraged for near term and long term operational benefit.

Ultimately we anticipate the effective use of KRR could result in:

- Quicker decision making to stay within a potential adversary's OODA loop;
- More robust AI systems capable of dealing with new information and handling uncertainty;

- Transparent system that give sufficiently understandable and assessable output;
- The retention of subject matter expertise as staff rotate through operations or as operations close, but then allowing the more rapid stand up of previous capabilities; and
- Greatly expanded ability to leverage and discern knowledge from existing data holdings.

1.5 APPROACH AND THE STRUCTURE OF THIS REPORT

This report is aimed at the wider NATO STO community and national representatives who:

- May be required to lead technology change initiatives and may benefit from using KRR methods and approaches in military contexts;
- Need to implement new solutions to make better use of information and knowledge; and
- May have expertise in ML and be seeking additional approaches to increase the robustness and explainability of results.

Firstly, this report focuses on the core concepts of knowledge representation (Chapter 2), recognising the first step in exploiting knowledge-based methods is having the means to represent knowledge, before then moving to approaches for reasoning over knowledge, or in other words the methods to deduce new knowledge from that which we already know. The report then moves on to the issues of implementing KRR approaches (Chapter 3), using specific examples to illustrate the issues involved. Finally, we present a short discussion on active research themes (Chapter 4) prior to offering conclusions and recommendations (Chapter 5).

1.6 REFERENCES

[1] Simon, H.A. (1971). Designing Organizations for an Information-Rich World. In Computers, Communication, and the Public Interest, M. Greenberger (ed.). Baltimore, MD, John Hopkins Press, pp. 40-41.

[2] NATO (04 Dec 2019). London Declaration. Issued by the Heads of State and Government participating in the meeting of the North Atlantic Council in London 3 – 4 December 2019.

[3] NATO (21 Nov 2019). First NATO AGS Remotely Piloted Aircraft Ferries to Main Operating Base in Italy. https://www.nato.int/cps/en/natohq/news_171171.htm. Retrieved 12/11/2020, 2020.

[4] Allen, G. and Chan, T. (2017). Artificial Intelligence and National Security. Belfer Center for Science and International Affairs Harvard Kennedy School, Cambridge, MA.

[5] GOV.UK (30 Sep 2020). Chief of the Defence Staff, General Sir Nick Carter Launches the Integrated Operating Concept, Ministry of Defence.

[6] Ministry of Defence. (2018). Information Advantage.

[7] Office of the Director of National Intelligence (Jan 16 2019). The Aim Initative – A Strategy for Augmenting Intelligence Using Machines.

[8] Blasch, E., Steinberg, A., Das, S., Llinas, J., Chee, C., Kessler, O., Waltz E. and White, F. (2013). Revisiting the JDL Model for Information Exploitation, ISIF, Intl Society of Information Fusion, p. 129.

[9] NATO STO, (2019). 2019 Highlights Science and Technology Organization, NATO, p. 56.

[10] Pearl, J. and Mackenzie, D. (2018). The Book of Why, The New Science of Cause and Effect. Basic Books.

Chapter 2 – KNOWLEDGE REPRESENTATION AND REASONING – CORE CONCEPTS

2.1 INTRODUCTION

In the following section we discuss what is meant by Knowledge Representation and Reasoning (KRR). We begin by introducing the concept of knowledge and knowledge systems. We conclude this section on the relationship between the 'old-school' AI of KRR and recent advances in ML (particularly deep learning).

2.2 DEFINING KNOWLEDGE AND KNOWLEDGE SYSTEMS

Knowledge is one of those words that everyone knows the meaning of, yet finds hard to define. Other words such as data, facts and information are often used interchangeably with knowledge. However, the majority of academics and knowledge management authorities make a distinction between these related, but discrete terms [1] (see Figure 2-1).

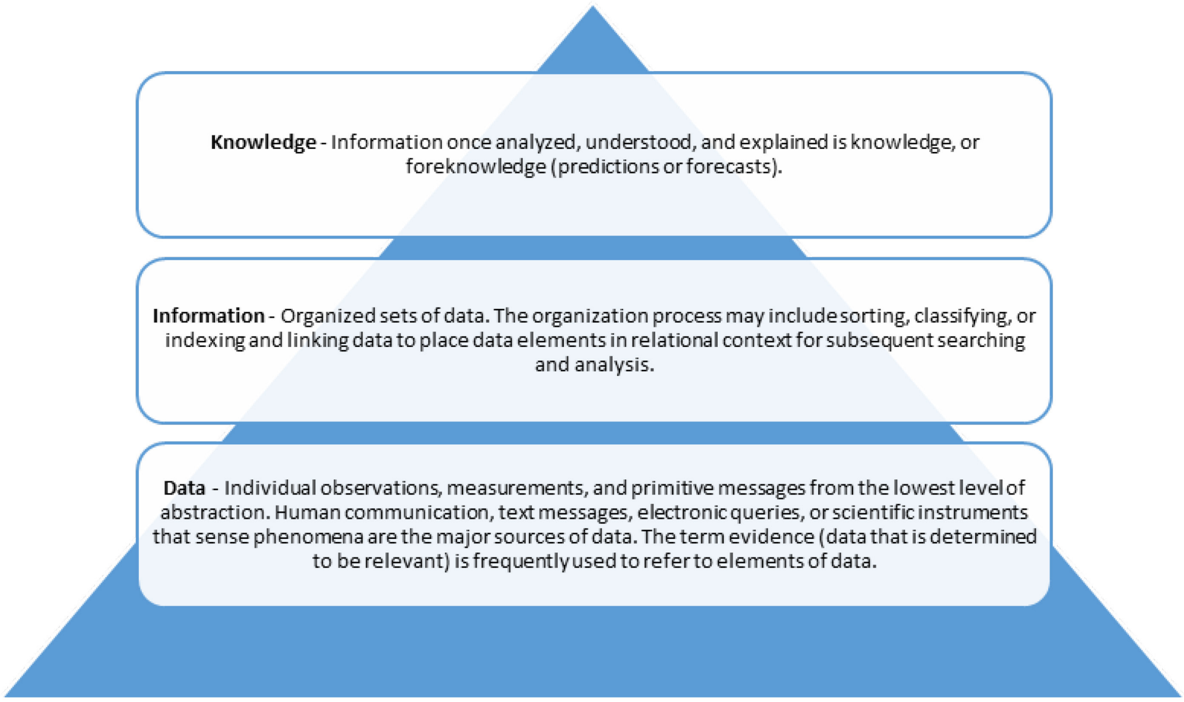

Figure 2-1: Distinguishing Knowledge, Information and Data (Adapted from Ref. [2]).

In creating knowledge the act of understanding information provides:

1) A degree of comprehension of both the static and dynamic relationships of the objects represented by the data;
2) The ability to model structures; and
3) Past (and future) behaviour of those objects.

Knowledge includes both static content and dynamic processes.

Knowledge, as the word is used for knowledge systems, refers to the codified experience of agents. Codified emphasises that knowledge is written (recorded). Experience emphasises that knowledge is created and used in experiential situations. Agents undergo experiences.

The term *knowledge system* is a shorthand for the term *knowledge-based system* [3]. A knowledge system is a computer system that represents and uses knowledge to carry out a task. An *expert system* is an intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution [4]. As the applications for the technology have broadened, the more general term *knowledge system* has become preferred by some people over *expert system* because it focuses attention on the knowledge that the systems carry, rather than on the question of whether or not such knowledge constitutes expertise. Figure 2-2 illustrates the basic concept of a knowledge-based (expert) system.

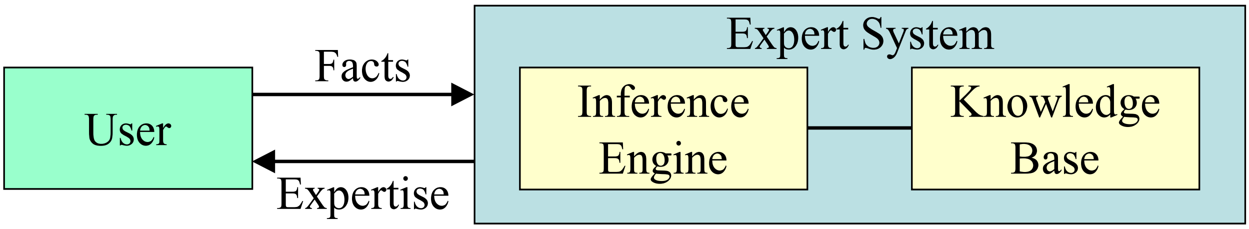

Figure 2-2: Basic Concept of a Knowledge-Based (Expert) System Function.

A knowledge-based system has by design the ability to be told facts about its world and adjust its behaviour correspondingly, thus the advantages of knowledge-based systems include that one can [5]:

- Add new tasks and easily make them depend on previous knowledge;
- Extend the existing behaviour by adding new beliefs;
- Debug faulty behaviour by locating the erroneous beliefs of the system; and,
- Concisely explain and justify the behaviour of the system.

A knowledge base is the organised repository for the collection of knowledge related to a domain and used for understanding, formulating, and solving problems in a knowledge-based system. Once a knowledge base is built, AI techniques are used to give the computer an inference capability based on the facts and relationships contained in the knowledge base. That is, the knowledge base contains a data structure that can be manipulated by an inference system that uses search and pattern matching techniques on the knowledge base to answer questions, draw conclusions, or otherwise perform an intelligent function.

2.3 EXPERTISE AND KNOWLEDGE ENGINEERING

Expertise is a specialised type of knowledge that is known only to a few [6]. It is not commonly found in public sources such as books and papers. In general, the term expert connotes both specialisation in narrow problem-solving areas or tasks and substantial competence [5], [7]. An expert's knowledge is specific to one problem domain, as opposed to knowledge about general problem-solving techniques. Expertise in one problem domain does not automatically carry over to another. Expert systems are generally designed to be experts in one problem domain. In fact, restricting the problem domain is typically necessary to produce useful solutions. Experts heavily rely on a vast knowledge of heuristics and the experience they have built up over the years. If an expert cannot solve a problem based on expertise, then he or she must reason from the first principles and theory.

The process of building a knowledge base is called knowledge engineering [1]. It deals with knowledge acquisition, knowledge representation, knowledge validation, inferencing, explanation and justification, and maintenance. Figure 2-3 offers a diagrammatic view. A knowledge engineer is someone who investigates a particular domain, determines what concepts are important in that domain, and creates a formal representation of the objects and relations in the domain [1]. A knowledge engineer is generally not an expert in the domain at hand. Knowledge acquisition is the process by which a knowledge engineer collects, extracts, transfers, accumulates, structures, transforms and organises knowledge (e.g., problem-solving expertise) from one or more knowledge sources (human experts, books, documents, sensors, or computer files) for constructing or expanding a knowledge base [5].

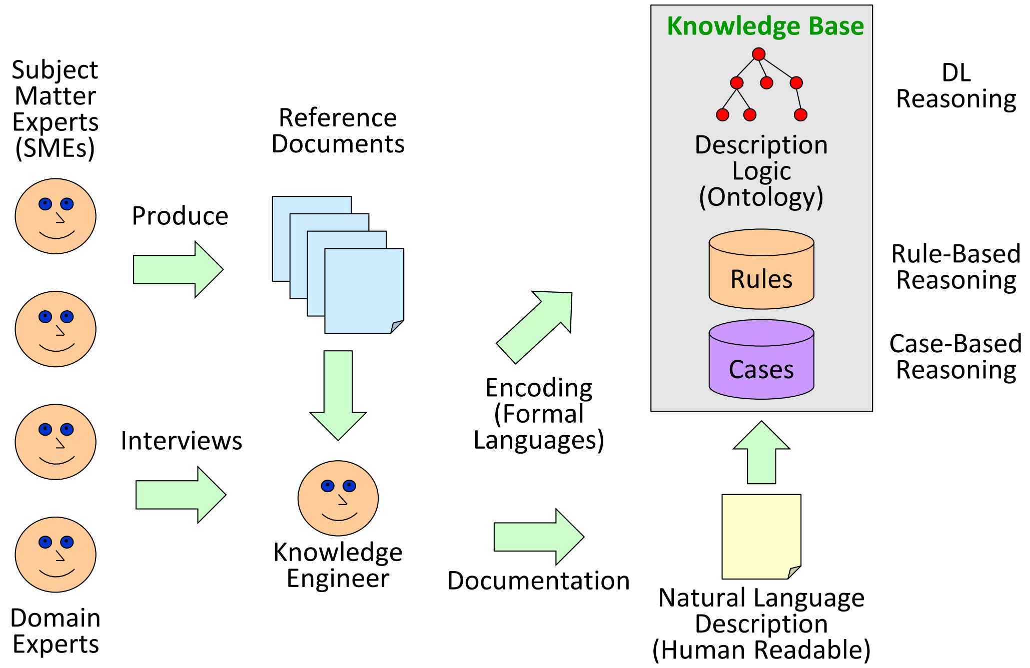

Figure 2-3: The Knowledge Engineering Process.

Unfortunately, acquiring knowledge from experts is a complex task. This process has been identified by many researchers and practitioners as a bottleneck that currently constrains the development and construction of expert systems and other AI systems. The process of representing knowledge of a domain goes through several stages. The following five-step methodology can be used [1]:

1) Decide what to talk about;
2) Decide on a vocabulary of predicates, functions, and constants;
3) Encode general knowledge about the domain;
4) Encode a description of the specific problem instance; and
5) Pose queries to the inference procedure and get answers.

One cannot put the world in a computer, so all reasoning mechanisms must operate on representations of facts, rather than on the facts themselves [1]. The object of knowledge representation is to express knowledge in computer-tractable form, such that it can be exploited. Knowledge representation research studies the problem of finding a language in which to encode knowledge so that the machine can use it [8]. It should support the tasks of acquiring and retrieving knowledge, as well as subsequent reasoning [5].

A language is a set of expressions and a set of combinatory rules. There are many types of languages; of particular interest are knowledge representation languages. A good knowledge representation language should combine the advantages of natural languages and formal languages [1]; it should be expressive and concise, be unambiguous and independent of context. A language is defined by two aspects: syntax and semantics [1]; the syntax describes the possible configurations that can constitute sentences, while semantics is an approach for assigning meanings to symbols and expressions. From the syntax and semantics, an inference mechanism can be derived that uses the language. A representation is formal when its symbols are interpretable by a computer program that uses them to guide its activity in carrying out a task.

One may have the impression that a knowledge engineer must find a single best representation and stick with it. However, it is not necessary to select and use a single representation in knowledge systems. Actually, no single knowledge representation method is ideally suited by itself for all tasks [5]. An important alternative is the use of multiple representations. A variety of knowledge representation paradigms, schemes and techniques have been devised over the years. These includes ontologies, knowledge graphs, lists and outlines, decision tables, decision trees, state and problem spaces, production rules, object-attribute-value triples, semantic networks, schemata, frames, scripts, logics, etc.

In the knowledge representation domain, a lot of attention has been devoted to ontologies and ontological engineering [9]. The word ontology was taken from Philosophy, where it broadly means a systematic explanation of being. The definition: "An ontology is a formal, explicit specification of a shared conceptualization" (Studer, Benjamins et al. [10] and based on the work of Thomas [11]) seems appropriate for the development of knowledge-based systems.

The ontology community distinguishes lightweight and heavyweight ontologies [9]; lightweight ontologies include concepts, concept taxonomies, relationships between concepts, and properties that describe concepts, while heavyweight ontologies add axioms and constraints. Finally, different knowledge representation techniques can be applied to model ontologies, although not all of them can represent the same knowledge with the same degree of formality and granularity. One of the key decisions to take in the ontology development process is to select the language (or set of languages) in which the ontology will be implemented. One benefit of using current W3C standards for ontology development is that the modelling conventions allow for the co-existence of synonym terms and multiple definitions for terms so that different domains of knowledge or different communities of experts can be more closely aligned through the model.

2.4 INFERENCING AND REASONING

In this section, we turn from the ways we can create formal representations of knowledge to the ways we can reason with our knowledge and related information that has been formally organised. We start with some brief definitions and then move to talk about the various approaches.

To reason is:

1) To use the faculty of reason so as to arrive at conclusions; or
2) To discover, formulate, or conclude by the use of reason [12].

Similarly, to infer is to derive a conclusion from facts or premises. The terms *reasoning* and *inference* are generally used to cover any process by which conclusions are reached. The term *inference* is generally used for mechanical systems such as expert systems, while *reasoning* is generally used in human thinking.

Once the knowledge representation in a knowledge base is completed, or is at least at a sufficiently high level of accuracy, it is ready to be used for reasoning tasks. One needs a computer program to access the knowledge for making inferences. This program is usually called the inference engine or the control program. It is used to direct the search through the knowledge base and to control the reasoning process.

Two general methods of inferencing are commonly used as problem-solving strategies: forward chaining and backward chaining. Forward chaining (also known as bottom up reasoning) is reasoning from facts to the conclusion(s) resulting from those facts. Its inference processes are not directed toward solving any particular problem; for this reason it is also called a data-driven or data-directed procedure. Backward chaining starts with something one wants to prove, find implication sentences that would allow him/her to conclude it, and then attempt to establish their premises in turn. Backward chaining thus involves "reasoning in reverse". It is normally used when there is a goal to be proved. It is also called a goal-driven or goal-directed procedure, in which one starts from an expectation (hypothesis), then seek evidence that supports (or contradicts) the expectation. Reasoning from the higher-level constructs such as hypotheses down to the lower-level facts which may support the hypotheses is called top-down reasoning. Whether forward or backward chaining is better depends on the purpose of the reasoning and the shape of the search space. For example, if the goal is to discover all that can be deduced from a given set of facts, the system should run forward. In some cases, the two strategies can be mixed (i.e., bidirectional chaining).

The standard patterns of in inferencing are based on philosophical underpinnings defined over 2000 years ago. *Modus ponens* (Latin for mode that by affirming affirms) and *modus tollens* (Latin for "mode that by denying denies"), among other patterns, offer rules that provide the basic building blocks of propositional logic. Based on these rules, and many others, a variety of approaches have been devised and used to achieve reasoning/inference in computer-based systems. Figure 2-4 describes typical approaches.

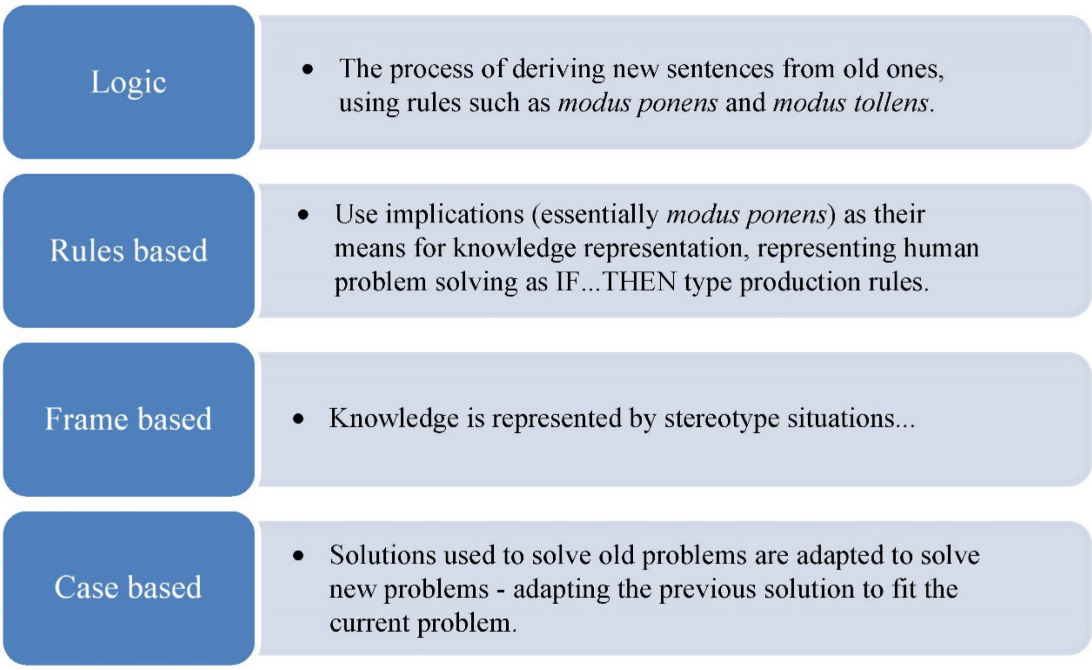

Figure 2-4: Representation Approaches for Reasoning and Inferencing.

2.5 KNOWLEDGE GRAPHS

Straddling graph theory and web science, the concept of the knowledge graph has emerged as a means of bringing together known 'facts' (or axioms) regarding some domain of discourse. Knowledge graphs have become the most important denominator in successful search engines on the World Wide Web. Social network sites (e.g., Facebook, etc.) and e-commerce sites (e.g., Amazon, etc.) are also using knowledge graphs to store and retrieve useful information.

There is no well-established definition of knowledge graphs, further discussion can be found in Ref. [13]. The following might be considered key characteristics for a knowledge graph in the context of ET-111:

- That information and knowledge can be captured in a structured representation without the need for a rigid, proscriptive schema;
- Where a formal, explicit descriptive model of a domain is required, the structure of the knowledge graph can be specified using formal ontology, which gives robust semantic descriptions of the classes and relationships within the domain of discourse; and
- That individual entities (individuals or instances) within a knowledge graph can then be described in terms of the classes and relationships defined by an ontology.

Knowledge graphs may also be considered examples of semantic networks, which when applied and combined to the World Wide Web is described as the Semantic Web [14].

2.6 SEMANTIC ENABLEMENT AND INTEROPERABILITY

A key application for ontologies is to provide the semantic 'glue' to enable systems to interoperate, to be able to exchange content with no loss of meaning. The principle of interoperability is not new and has for many years been accomplished through the use of mappings between one schema and another. So to transfer information from System1 to System2, a mapping between the schema of System1 and that of System 2 would be produced. This mapping enables exchange to be undertaken in which the content from one schema element can be transferred to another schema element. Such direct mappings between schemas can become numerous when multiple systems are involved, hence the use of hub and spoke models in which a third schema is introduced to act as an intermediary. The use of an intermediary reduces the number of schema mappings required from *n!* to *n* where n is the number of systems.

Whilst the use of a hub and spoke model for exchange of information is beneficial, a schema based approach focuses on the structure of the information rather than the content; it is not necessarily the case that information can be exchanged without loss of meaning. This is known as syntactic interoperability. The use of an ontology as the hub for exchange of information allows the semantics as well as the syntax of the information to be described, thus enabling semantic interoperability in which information can be exchanged without loss of meaning. The process then becomes one of mapping from an internal schema to some external ontology which fulfils the role of hub. All systems can map to the ontology and use import/export routines to exchange information in a semantically rich and platform agnostic manner. A number of ontologies are used in this space including Basic Formal Ontology [15] and its domain specific extensions. Recent work in the UK has taken a schema based model and used it as the basis for an ontology with the specific aim of providing an ontology driven UK Intelligence Community (UKIC) Information Exchange Standard (IES), described in more detail in the use cases section (Figure 2-5).

Associated with the concept of interoperability is that of enablement. It is often the case that information systems make use of internal schema optimised to their purpose. It may not be possible or desirable to modify the internal schema of any information system but semantic interoperability is required above and beyond import/export which can be achieved by mapping the schema to the ontology. In which case, there is a need to put an additional layer on top of the information system to achieve this. This additional semantic enablement layer mediates between the external ontology and the internal schema and can be achieved in a variety of ways depending on the use case and the style of architecture in use. One such approach is Ontology Based Data Access (OBDA) which leverages an ontology and a mapping to the source schema, as exemplified by the Ontop framework [16], [17], thereby supporting semantically mediated interaction with an underlying information repository. In the geospatial community, a service driven approach has been developed building on existing Open Geospatial Consortium (OGC) standards for Spatial Data Infrastructures (SDI) of the kind widely in use, including within UK MOD. This approach has focussed on profiling existing geospatial web services to provide a semantic enablement layer [18]). Similar approaches have also been explored in domains including Internet of Things (IoT) [19] and Sensor Web Enablement [20].

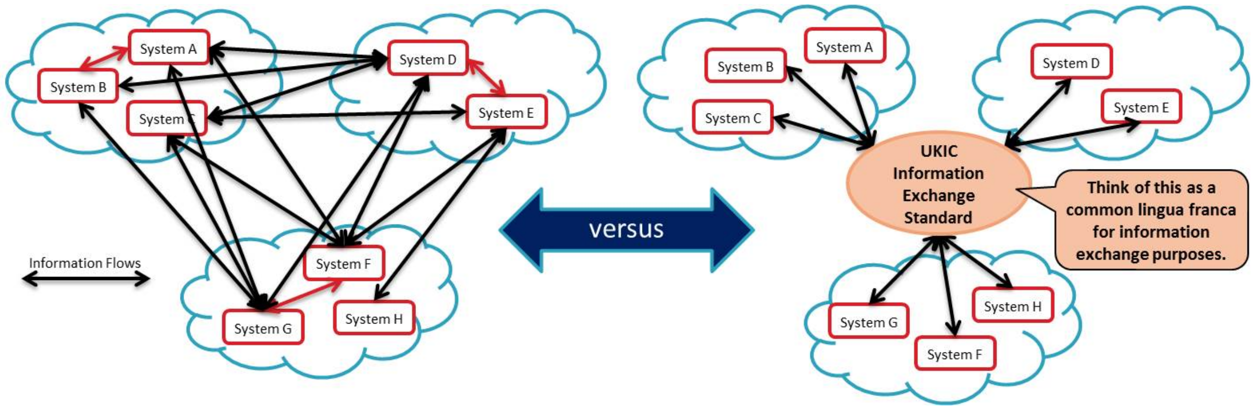

Figure 2-5: Use of the UKIC IES for Exchange of Information.

2.7 UNCERTAINTY MANAGEMENT

In this section, we provide a high level overview of the fundamentals of uncertainty management, in the context of KRR. Further detail is provided in Annex E.

Most tasks requiring intelligent behaviour have some degree of uncertainty associated with them. The type of uncertainty that can occur in knowledge-based systems may be caused by problems with the data. For example [21]:

- Data might be missing or unavailable;
- Data might be present but unreliable or ambiguous due to measurement errors;
- The representation of the data may be imprecise or inconsistent;
- Data may just be user's best guess; and
- Data may be based on defaults and the defaults may have exceptions.

The uncertainty may also be caused by the represented knowledge, since it might:

- Represent best guesses of the experts that are based on plausible or statistical associations they have observed; and/or
- Not be appropriate in all situations (e.g., it may have indeterminate applicability).

Given such numerous sources of errors, most knowledge-based systems require the incorporation of some form of uncertainty management.

When implementing an uncertainty scheme, one must be concerned with three issues:

- How to represent uncertain data;
- How to combine two or more pieces of uncertain data; and
- How to draw inference using uncertain data.

A review of the main typologies proposed in the literature [22] considers a number of problems regarding the different types of uncertainty, the different epistemic interpretations, the different mathematical representations, in order to better understand and use the existing mathematical formalisms for reasoning under uncertainty.

The term "uncertainty reasoning" is meant to denote the full range of methods designed for representing and reasoning with knowledge when Boolean truth values are unknown, unknowable, or inapplicable [23]. To illustrate, consider a few reasoning challenges that could be addressed by reasoning under uncertainty:

- Automated agents are used to exchange information that in many cases is not perfect. Thus, a standardised format for representing uncertainty would allow agents receiving imperfect information to interpret it in the same way as was intended by the sending agents.
- Much information is likely to be uncertain. Examples include weather forecasts or gambling odds. Canonical methods for representing and integrating such information are necessary for communicating it in a seamless fashion.
- Information is also often incorrect or only partially correct, raising issues related to trust or credibility. Uncertainty representation and reasoning helps to resolve tension amongst information sources having different confidence and trust levels.
- Many visions rely on numerous distinct but conceptually overlapping ontologies that co-exist and interoperate. It is likely that in such scenarios, ontology mapping will benefit from the ability to represent degrees of membership and/or likelihoods of membership in categories of a target ontology, given information about class membership in the source ontology.
- Dynamic composability of services requires runtime identification of processing and data resources and resolution of policy objectives. For some such cases, uncertainty reasoning techniques may be necessary to resolve situations in which existing information is not definitive.
- Information extracted from large information networks is typically incomplete. The ability to exploit partial information is very useful for identifying sources of service or information. It is clear that search effectiveness could be improved by appropriate use of technologies for handling uncertainty.

As work with semantics and services grows more ambitious, there is increasing appreciation of the need for principled approaches to representing and reasoning under uncertainty [24].

To model uncertainty, many mathematical tools have been developed, being either qualitative such as modal or nonmonotonic logics, or quantitative approaches such as probability theory, fuzzy sets theory, or evidential theory [22]. These approaches are often compared on the basis of their different strengths and weaknesses:

- Their suitability to model a particular type of uncertainty;
- Their requirement for prior knowledge;
- Their computational time complexity;
- The need for independence constraints; and
- Their reasoning capacities.

2.8 SYMBOLIC VERSUS SUB-SYMBOLIC APPROACHES

Within AI, we distinguish symbolic from sub-symbolic approaches. Sub-symbolic approaches comprise in particular ML techniques, while symbolic approaches apply the techniques outlined above, like, e.g., knowledge graphs.

We distinguish four classes of sub-symbolic ML approaches, namely supervised learning, unsupervised learning, reinforcement learning and hybrid approaches. All of these require a training phase, in which models are created from training data or parameters of given but underspecified models are optimised with regard to a specific performance criterion. The trained models serve the structuring and understanding of the given data (e.g., clustering), predictions based on newly received data (e.g., classification, statistic forecast by regression analysis), the target-oriented creation of data (e.g., text or image generation), or the derivation of a strategy for performing optimal sequences of actions that lead to a desired result. A variety of methodologies for implementing these approaches are on the market, ranging from linear and logistic regression, via Bayesian networks, decision trees, support vector machines and random forests, to (deep) neural networks, just to name a few. All methodologies require the initial specification of the features for which optimisation is to be achieved through training. As the manual specification of these features is laborious and can become a bottleneck in the creation of ML systems, feature or representation learning approaches are being developed for the further automation of this process. Moreover, to change an ML model, it usually has to be trained from scratch again. Even for only slight changes this can require enormous effort. Therefore, approaches to the adjustment of trained models, e.g., through lifelong ML [25] or transfer learning, are ongoing research challenges.

Sub-symbolic ML systems are black boxes that make it hard or even impossible for a user to understand how and why a certain output has been generated. It is thus at least difficult for a user to validate the output of such a system and become aware of potential malfunctions. It is obvious, that (necessarily) overtrusting users are vulnerable to attacks on their systems [26]. It is, therefore, a challenge to create understandable systems.

With every passing year, sub-symbolic ML and knowledge representation learning on knowledge graphs – that is, learning of symbolic representations [27]) – are advancing rapidly, both in scale and depth, but in different directions. On the one hand, ML techniques are getting better at performing various tasks (classification, generation, etc.) on a variety of datasets with great precision and recall. On the other hand, knowledge representation brings the ability to represent entities and relations with high reliability, explainability and reusability.

Symbolic approaches are best suited for manageable and completely defined application scenarios. Within such scenarios, they can enable control and automation of (standard) processes. For the design of a symbolic system, both the problem and its solution must be well understood. (ML systems, on the contrary, can be designed without a prior full understanding of the solution – it is sufficient to understand the problem and the database.) Quite often we do not have a choice: if data for training an ML-system are not available, we must refer to expert knowledge and either create a symbolic system or create a rule-based data generator which creates the input for an ML-system.

The integration of symbolic and sub-symbolic systems as well as the use of symbolic systems as subsidiary systems for ML systems can systematically improve the accuracy of the systems and extend the range of ML capabilities. In particular, results inferred from ML models can be given better explainability and trustworthiness.

2.9 SUMMARY

This section has introduced some of the key concepts underpinning KRR – defining the concept of knowledge and the describing the structure of a knowledge system. It has discussed, and given examples of knowledge representation, emphasising the importance of domain expertise as the process of knowledge engineering seeks to build the underpinning knowledge base. The use of inferencing over that knowledge base has been explored and the variety of methods available to do this have been introduced. Finally, it has discussed the recent introduction of knowledge graphs, and compared such symbolic methods with those represented by ML.

Key takeaways from this section include:

- Knowledge is created and used in experimental situations, and must be recorded (written down) in order to exist. Knowledge-based systems provide the means to represent and interact with knowledge.
- The process of acquiring, representing, verifying and managing knowledge is a complex task and has previously constrained the development of such knowledge-based system.
- Not all knowledge can exist in a computer, approximations must be made, meaning any reasoning over this knowledge must operate on representations of facts, not the facts themselves. The representation of those facts requires an expressive yet concise approach that is unambiguous and independent of context. It is not necessary to have only one method of representation.
- Ontologies and ontological engineering are key tools in knowledge representation. They are generally distinguished between so called lightweight (definition of concepts and relationships between concepts) and heavyweight (the addition of axioms and constraints) ontologies.
- The term inferencing is generally reserved to the activities of a machine in reaching a conclusion based on available data, whereas reasoning generally refers to human thinking.
- Any representation requires some approach to uncertainty management.
- Knowledge graphs, building on representation methods offer symbolic approaches for AI and contrast with ML. Whereas ML generally requires limited understanding of the domain (just of the data and the problem statement) symbolic methods rely heavily on domain expertise. However, in some cases if training data is not available (for example when considering rare events) we must refer to symbolic systems built on human expertise.

Finally, we note that knowledge representation is a prerequisite for reasoning (inferencing), but knowledge representation can be done for other purposes such as improved search and discovery of information.

2.10 REFERENCES

[1] Girard, J.L. (2004). Defence Knowledge Management: A Passing Fad? Canadian Military Journal, Summer 2004, pp. 17-27.

[2] Rowley, J. (2007). The Wisdom Hierarchy: Representations of the DIKW Hierarchy. Journal of Information and Communication Science 33(2), pp. 163-180.

[3] Stefik, M. (2014). Introduction to Knowledge Systems. Morgan Kaufmann.

[4] Giarratano, J. (2005). Expert Systems: Principles and Programming. Thomson, Switzerland.

[5] Brachman, R. and Levesque, H. (2004). Knowledge Representation and Reasoning. Morgan Kaufmann.

[6] Nonaka, I. (July – August, 2007). The Knowledge-Creating Company. Harvard Business Review.

[7] Webster’s (1981). Webster’s Third New International Dictionary. Springfield MA, Merriam Webster.

[8] Ginsberg, M. (1993). Essentials of Artificial Intelligence, Morgan Kaufmann.

[9] Gómez-Pérez, A., Fernández-López, M. and Corcho, O. (2004). Ontological Engineering, Springer.

[10] Studer, R., Benjamins, R. and Fensel, D. (2007). Knowledge Engineering: Principles and Methods. Germany, Europe, Universität Karlsruhe.

[11] Thomas, R.G. (1993). A Translation Approach to Portable Ontology Specifications.

[12] 11th Collegiate Dictionary, Merriam-Webster. (2003).

[13] Ehrlinger, L. and Wöß, W. (2016). Towards a Definition of Knowledge Graphs. 1st International Workshop on Semantic Change & Evolving Semantics (SuCCESS16). Leipzig, Germany, p. 1695.

[14] Bonatti, A., Decker, S., Polleres, A. and Presutti, V. Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web, Dagstuhl Seminar 18371, pp. 89-92.

[15] Smith, B., Almeida, M., Bona, J., Brochhausen, M., Ceusters, W., Courtot, M., Dipert, R., Goldfain, A., Grenon, P. and Hastings, J. (2014). Basic Formal Ontology 2.0 Draft Specification and User's Guide.

[16] Rodriguez-Muro, M., Kontchakov, R. and Zakharyaschev, M. (2013). Ontology-Based Data Access: Ontop of Databases. International Semantic Web Conference, pp. 558-573.

[17] Bagosi, T., Calvanese, D., Hardi, J., Komla-Ebri, S., Lanti, D., Rezk, M., Rodriguez-Muro, M., Slusnys, M. and G. Xiao (2014). The Ontop Framework for Ontology Based Data Access. Chinese Semantic Web and Web Science Conference, pp. 67-77.

[18] Janowicz, K., Schade, S., Bröring, A., Keßler, C., Maué, P. and Stasch, C. (2010). Semantic Enablement for Spatial Data Infrastructures. Transactions in GIS 14(2), pp. 111-129.

[19] Gilani, K., Kim, J., Song, J., Seed, D. and Wang, C. (2018). Semantic Enablement in IoT Service Layers – Standard Progress and Challenges. IEEE Internet Computing 22(4), pp. 56-63.

[20] Bröring, A., Echterhoff, J., Jirka, S., Simonis, I., Everding, T., Stasch, C., Liang, S. and Lemmens, R. (2011). New Generation Sensor Web Enablement. Sensors 11(3), pp. 2652-2699.

[21] University of Illinois. (2020). Chapter 4. Reasoning Under Uncertainty. Retrieved 01 May 2020, from https://www.cs.uic.edu/~liub/teach/cs511-spring-06/cs511-uncertainty.doc.

[22] Jousselme, A.L., Maupin, P. and Bosse, E. (2003). Uncertainty in a Situation Analysis Perspective, IEEE. 2, pp. 1207-1214.

[23] Laskey, K., Laskey, K., Costa, P., Kokar, M., Martin T. and Lukasiewicz, T. (2008). Uncertainty Reasoning for the World Wide Web: Report on the URW3-XG Incubator Group. United States, North America.

[24] Costa, P.C.G., Laskey, K.B., Blasch, E. and Jousselme, A.L. (2012). Towards Unbiased Evaluation of Uncertainty Reasoning: The URREF Ontology, IEEE, pp. 2301-2308.

[25] Chen, Z. and Liu, B. (2018). Lifelong Machine Learning, Morgan & Claypool.

[26] Vorobeychik, Y. and Kantarcioglu, M. (2018). Adversarial Machine Learning, Morgan & Claypool.

[27] Hamilton, W. (2020). Graph Representation Learning, Morgan & Claypool.

Chapter 3 – IMPLEMENTING KNOWLEDGE REPRESENTATION AND REASONING

The following section presents a discussion on several of the engineering implications that must be kept in mind when using KRR techniques. This includes choices about the construction of knowledge models themselves, as well as ensuring that any tool is fit for purpose. A core consideration here is the need for data integration and interoperability between systems, and the ability for systems to be easily extended using capabilities developed elsewhere (and perhaps for other purposes). Therefore, the choice of standards and overarching architecture are important. Paradigmatic examples for the implementation of Knowledge-Based Systems (KBSs) – following different purposes – are given. We conclude this chapter with a discussion of common concerns as well as strengths and weaknesses regarding KBSs.

3.1 ONTOLOGIES FOR INTEGRATION, INTEROPERABILITY, AND INFORMATION SHARING

Knowledge representation, in particular information models and ontologies, can serve as Semantic Reference Models (SRMs) for data management and interoperability. Data and information management in the C2 domain demands the integration of distributed information from various sources and heterogeneous systems, including legacy systems. This information can be stored in various data silos or in data lakes that comprise heterogeneous formats and standards. To properly integrate the information, translate one format into another and derive actionable information by combining various sources, an SRM is needed by which a common semantics for the interpretation of the available data is specified. The concept of such an SRM is, thus, part of the NATO Core Data Framework. (So far, however, it is just a concept.)

The challenge of information integration becomes harder when different partners with their systems come into play. Then, a proper interoperable solution is required. Interoperability is the capability to exchange messages between at least two partners, to read and interpret these messages in a consistent way – that is, the partners must have a mutual understanding of the information they exchange – and, finally, to react on the messages in a foreseeable manner. Interoperability involves at least two sides where operational subject matter experts perform operational activities that are expected to be consistent, coordinated, and contributing to a common goal. These activities mostly rely on information systems – in the military domain: C2 Information Systems (C2IS) – that process data, some of which can be exchanged between systems. Thus, an interoperability solution implies human operators, information systems including local processing services and exchange services and, finally, information to be exchanged (implying that the sharing of information has been agreed upon).

Ontologies and information models are representations of concepts, their attributes and relations. They represent chosen subject areas, i.e., domains of discourse. One might technically distinguish ontologies from information models: information models mostly follow an object-oriented approach with the Unified Modelling Language (UML) as a representation format, while ontologies are rather based on the Web Ontology Language (OWL) or some other graph representation format. Here, we regard the terms “ontology” and “information model” as synonyms.

Within the last decades, a plethora of standards and reference models have been created within the military domain as well as the intelligence domain, for various subdomains and purposes, from intra-vehicle data exchange via coordination of human, robotic forces and simulation systems to headquarter-to-headquarter communication. In principle, every NATO STANAG comes with the fragment of an ontology, e.g., a taxonomy of ships (APP-20) or fuels (Logistics Handbook). Closest to the common understanding of an ontology are, most probably, the C2SIM (C2 Simulation Interoperation) ontology [1], [2], the MIM, and the

Defence Intelligence Core Ontology (DICO). The C2SIM ontology is based on prior work on the Coalition Battle Management Language C-BML [3]. It serves as a semantic reference for communication with simulation systems and has also been applied for information exchange with unmanned systems.

As a means for standardizing messages and thus enabling interoperability, the US National Information Exchange Model (NIEM) is frequently mentioned, too [4]. The NIEM, however, cannot be considered a proper ontology. It provides a vocabulary which Communities of Interest refer to in order to define Information Exchange Package Documentations (IEPDs). IEPDs are, in essence, message text formats. The NIEM is missing the expressive means to define concepts together with their properties and relations. It is, thus, closer to a standard of message text formats like NATO APP-11 than to an ontology.

The MIP Information Model (MIM)

The MIP (Multilateral Interoperability Programme) Information Model defines common semantics for the Command & Control (C2) domain. The main objective of the MIM is to support information exchange in joint and combined operations.

- The core of the MIM is a taxonomy with thousands of militarily relevant concepts.
- Concepts in the MIM have a rich set of properties to describe the characteristics of – and relationships between – their instances.
- The MIM is platform independent, so it is not tied to a specific exchange technology.
- Communities of Interest (COIs) can adopt the MIM for developing interoperability specifications in support of their specific processes. The model can be extended, or have subsets added.

Refer to the MIM entry in the annexes for further detail.

Rich Event Ontology (REO)

The Rich Event Ontology (REO) aims to represent lexical event semantic information; thus, it hopes to capture both common-sense world knowledge about events and their participants, as well as lexical information on how these concepts are realised and tagged in various English annotation schemas. REO:

- Unifies existing Semantic Role Labelling (SRL) schemas used in Natural Language Processing (NLP) by providing an independent conceptual backbone through which they can be associated.
- Augments the schemas with event-to-event causal and temporal relations.
- Facilitates reasoning on and across documents, revealing relationships between events that come together in temporal and causal chains.
- Part of the world model for autonomous agents in search and navigation experiments.

Refer to the REO entry in the annexes for further detail.

Dstl Underpinning Data Science (UDS) OPIS Ontology

OPIS is a Polish word meaning description. The approach to ontology being developed within the UDS project at Dstl is based on description of phenomena. This phenomenological approach starts with the premise that the world can be viewed in terms of phenomena which we describe in order to come to an understanding of the world around us. These phenomena comprise events (things which happen, bounded in space and time) and states (temporally bound conditions of things) which may be changed by events. The concepts of Structured Observation Management (SOM) and Object Based Production (OBP) have influenced this design in that observations (being descriptions of some phenomena) can be seen as structured observations whilst the objects involved in phenomena (both physical and conceptual) can be seen as objects in an OBP sense.

In practice, this results in an event-centric model in which information is the product of activities (a form of event with some output) which produce descriptions, or observations, of real-world events; these activities being carried out by agents, i.e., humans or machines. Furthermore, properties, characteristics and relationships are modelled as state (a kind of phenomena) making their temporal nature explicit.

Provenance of observations is explicit and related work incorporates observations into formal models of argumentation as evidence/information nodes, thereby supporting formulation and evaluation of hypotheses. Importantly, this observation based approach directly supports uncertainty and multivocality as axioms within a knowledge graph using this ontology do not represent ‘truth’ rather they represent some proposition of a perceived ‘truth’, accepting that there may be many such perspectives; for example, when measuring the length of an object, there is undeniably some true length but due to the measuring process, all measuring activities can only ever approximate this within the constraints of method or instrument used (i.e,, precision) and nature of the specific measuring event (i.e., accuracy, fallibility, etc); similarly, multiple witnesses may claim an object as being of different colours or a car having different license plate number, etc.

The model is modular based on a tree analogy with a core providing a ‘trunk’ comprising specification of essential concepts such as Place, Actor, Physical and Conceptual Thing common to all applications. This is extended to the ‘branches’ through domain specific extensions, for example to cover maritime situational awareness and pharmaceutical production facilities (the two use cases conducted to date). Finally, detail is added at the ‘leaves’ through reference data used to enrich descriptions by means of standardised vocabularies such as those produced by UK MOD, US DOD and NATO.

This modular approach includes the use of external ontologies where appropriate. The model leverages extant standards for space (OGC GeoSPARQL) and time (W3C/OGC OWL-Time) which provide relative and absolute positioning for phenomena in some coordinate space, spatial and/or temporal. Further external ontologies are used to support provenance (W3C PROV-O) and uncertainty.

The model is based on the W3C Semantic Web stack with the ontology encoded using RDFS and OWL in order to support Description Logic based reasoning across the full extent of the model. SKOS is used to provide a lightweight structured vocabulary component for terminologies, allowing for a parallel governance regime to be applied to parts of the model likely to change.

The model makes extensive use of Ontology Design Patterns whereby a simple pattern is defined and then reused. For example, the act of making an observation uses a production pattern, as does the manufacture of a physical object or the creation of a design or plan.

Defense Intelligence Core Ontology (DICO)

DICO is a mid-level ontology, designed and built according to Basic Formal Ontology (BFO) standards [5]. The DICO is the Defense Intelligence All-source Analysis Enterprise (DIAAE) knowledge model for Object Based Production (OBP). It provides the semantic framework to access and organise defence intelligence data in a way that is intuitive and mission focused for the DIAAE analysts and collectors preparing for, and participating in, dynamic conflicts. Within the DICO, concepts are modelled using real world relationships that are structured in a way that is meaningful to both computers and humans.

This approach enables standards-based information exchange and interoperability between integrated applications and services. The expressivity and flexibility of the DICO knowledge model allows for future development, information sharing, and analytics. The DICO will facilitate the:

- Consistent development of classes and relationships that reflect the content found in authoritative Defense Intelligence Analysis Program (DIAP) sources such as the Modernized Integrated Database (MIDB);
- Ability to better incorporate spatio-temporal entities (e.g., the movement of mobile missiles out of garrison) with current and future analytic tool suites and databases focused on fixed entities such as facilities;
- Enhanced (i.e., computer-assisted tools such as machine learning) reasoning that supports intelligence analysis methods instead of data dictating analysis;
- Integration of relevant data from disparate intelligence sources and publicly available sources into a common object management service;
- Logical and consistent expansion of reasoning support into any domain at any level of granularity (i.e., from large aggregate objects down to elemental parts of objects); and the
- Improved usage of Intelligence Functional Codes (IFCs) and other Intelligence Community coding systems to reason with intelligence and analyse production.

The primary form of implementation of the DICO will be integration the Object Management Service (OMS) of the Defense Intelligence Agency (DIA). Creation and implementation of DICO compliant ontologies at the application level will strengthen semantic integration across the Defense Intelligence Enterprise and add to an evolving Enterprise Knowledge Graph.

Refer to the DICO entry in the annexes for further detail.

3.2 THE W3C SEMANTIC WEB STACK

The semantic web is an endeavour of the W3C to standardise the syntax and semantics of structured information that can be made available via the WWW, so that a web of (unambiguously interpretable) data is created. In addition to unstructured information comprising texts, images, video, etc., such structured information shall be accessible, integrate-able and process-able for arbitrary purposes. The Semantic Web comes with a set of open standards, called the W3C Semantic Web Stack [6].

An open standards approach provides a foundation that enables the use of standards compliant software, tools and architectural models, widely supported by industry and academia. The W3C Semantic Stack is a modern day foundation for an open standards approach for KRR. It is composed of the following standards:

- The eXtensible Mark-up Language (XML) as a mark-up language for defining syntax of information;
- The Resource Description Framework (RDF) and its schema (RDFS) for coding information triples which, in essence, denote relations between entities and, in combination, are to create a global knowledge graph;
- The OWL2 (the current version of OWL) as a description logic based language to define ontologies which goes beyond the expressiveness of RDFS. OWL supports inferencing and reasoning capabilities and is used by the OWL-Time ontology to describe temporal aspects and the Geospatial SPARQL Protocol and RDF Query Language (GeoSPARQL) ontology to describe geospatial aspects. The MIM is originally a UML-model but has been provided with a transformation to OWL;
- The Rule Interchange Format (RIF) and the Semantic Web Rule Language (SWRL) for defining and exchanging rules;
- SPARQL as a query language for RDF-triples; and
- The PROV ontology (PROV-O) which offers an approach to describe provenance.

The use of these open standards components allows the application of standards compliant approaches to reasoning, particularly the application of description logic using, for example, HermiT [7] or Pellet [8] reasoning engines. Military systems development can profit from W3C developments in two respects: firstly, it can profit from the structured information that is being made available via the semantic web. Secondly, it can technologically build upon the open standards and related technology.

3.3 CASE STUDIES

We offer the following examples of different cases where KRR is being used.

3.3.1 Building Domain Ontologies – DICO Development Process, Design Principles, and Best Practices

The Defense Intelligence Core Ontology (DICO) is a mid-level ontology, designed and built according to Basic Formal Ontology (BFO) standards [5]. A primary purpose of the DICO is to support the operators' – in this case intelligence analysts' – ability to reason with their data. Plain language questions can be fairly easily transformed into machine-readable query statements, to be sent to an analytic engine. The authors of the DICO captured concepts of common concern to defence all-source intelligence analysts. Further detail can be found in Annex C.

3.3.2 Knowledge Representation and Reasoning in Practice – The WISDOM R&D Platform

WISDOM is a R&D software platform [9] that has been developed at Defence Research and Development Canada (DRDC), mainly under Project 05da: Joint Intelligence Collection and Analysis Capability (JICAC), and is meant to be a proof-of-concept prototype of an intelligence production support system. It is geared towards research in data/information/knowledge integration, fusion, analytics, management and exploitation, aiming at providing a capability to support the analysts and decision makers in developing their belief, opinion, judgment, or prediction about situations while these individuals are involved in situation analysis and decision-making activities. A more detailed discussion of WISDOM can be found in Annex D.

3.3.3 Relevance Filtering, Information Aggregation and Enrichment – The Intelligent Situational Awareness (INSANE) Framework

Within the "Intelligent Situational Awareness" [10] framework at Fraunhofer FKIE, we conduct experiments in rule-based information management for C2IS. Operators of C2IS have to receive exactly all information that is relevant regarding their task and role, at the right time. They must neither be under-informed nor overtaxed. Therefore, C2IS has to provide suitable information management tools. These tools have information filters that allow for the creation of overlays in a Geographical Information System (GIS) or the selection of messages from a news source [11]. They assess the relevance of information with respect to the actual situation, on-going processes, the role and the current tasks of the operator.

One approach to the creation of information filters is rule-based: filtering criteria are manually defined as logical rules for assessing information items. The challenge is to make filters sufficiently fine-grained without running into the risk of making them overly complex, hard to understand and maintain, and eventually inconsistent. A possible solution is modularisation, that is, the definition of several coarse-grained filters and their appropriate combination. Alternatives to the rule-based approach can achieve a higher degree of automation by applying ML-techniques [12]. However, trained filters must be individually adaptable and, to that end, come in a human-understandable and -editable format. ML-approaches and rule-based approaches can be integrated by providing a user interface that enables the operator to give feedback and manually customise their profile. The challenge is in the proper (most probably: rule-based) orchestration of the different filters.

Filtering is only one part of information management. A system should also consider how information can be aggregated and enriched in such a way as to be useful to the operator. Let us assume that we are operators in the command post of a battalion. We are responsible for assessing the situation of the red forces (G2/S2). Firstly, we are provided with a map of the battle space. Within that map, relevant terrain properties are to be detected. To these belong potential borders between units, like rivers, heights, etc. Secondly, strategic reconnaissance provides us with an overview on the red forces to be expected on the battlefield. That is, to some extent we know which units we will be confronted with, how they will be organised and how they will be equipped. Thirdly, we have knowledge of the red doctrines – we know the space a tank platoon takes, and so on. Fourthly, we conduct reconnaissance ourselves and, thus, receive information, e.g., on sighting a tank.

We assume that a tank is always part of a company, either as a member of a platoon or as the tank of the company commander. From the sighting of a tank, we conclude that a company must be in that area. We know the space a company and its platoons take up and we know terrain properties that restrict the potential spaces. We can thus approximately locate the company. Sightings of additional Battle Space Objects (BSOs) enable us to further specify our situational overview. We aggregate the information we receive, close information gaps and determine the red disposition of forces, including the borders between units and the point of main effort. This can be achieved by a rule-based system: constraints imposed by terrain properties and doctrine are to be defined as rules. These rules plus facts derived from reconnaissance enable inferences about the actual situation.

Information aggregation and enrichment not only leads to a better situational overview by distilling and organising what we know but also by determining knowledge gaps and, therefore, directing further reconnaissance: if we have recognised two platoons next to each other, we know that there must be a third platoon but we might not know on which flank it is positioned. We must focus our reconnaissance accordingly. Also, if we recognise an unexpected vehicle, we should question our initial assumption on the red forces organisation and equipment.

One might be tempted to apply ML-techniques instead of symbolic knowledge representation in order to train an information management system. This, however, would require great amounts of realistic training and testing data which are most probably not available. Therefore, we have no choice except to create aggregation management support on the basis of symbolic knowledge representations.

3.3.4 Exchanging Information Within the UK and Five Eyes Defence and Security Community – The UKIC Information Exchange Standard (IES)

The Information Exchange Standard (IES) is a standard developed within the UK Intelligence Community by the UKIC Entity Working Group (EWG) for the purpose of exchanging processed/analysed intelligence. This is based around a particular kind of ontology, specifically a four-dimensional extensional ontology in which the criterion for identity of any particular entity is its spatio-temporal extent. This can be seen in the Figure 3-1, showing how states are used to record measurements, in this case the mass of a person over time.

Figure 3-1: Mass of a Person Over Time Using IES Representation.

This form of ontology in this example explicitly models a person in terms of their spatio-temporal extents, with mass being a characteristic of states of the person at different times.

The use of RDF Schema as the basis for this standard, combined with the mereological basis of its class hierarchy, provides a readily extensible yet semantically rich form of representation for the purposes of exchange (Figure 3-2).

3.4 OPPORTUNITIES AND CHALLENGES OF IMPLEMENTING KNOWLEDGE-BASED SYSTEMS

3.4.1 Discussion of Common Concerns

Let us discuss common concerns and objections against approaches to symbolic AI.

3.4.1.1 Should I Really Care, Don't ML Methods Deliver It All?

Question: KRR/symbolic methods are an established technology with a large body of knowledge behind them, and a suite of standards supporting the semantic web. However, the current focus in AI is on computational (ML) methods, largely driven by the convergence of data availability, processing power and software frameworks. Given the hype it could be interpreted that ML can deliver much of what knowledge-based (symbolic) methods could, without needing to go through the effort of defining ontologies. Is that true?

Reply: ML methods require data. The amount of data required to learn over a state space of *d* dimensions increases exponentially with *d*, so the data requirement can become unrealistic over a large-enough state space. In such cases where adequate data are not available, ML is not an option. A further important reason for considering knowledge-based methods is the need to ensure explainability. AI that learns independently and finds different patterns with each new set of data cannot be audited against an accepted cognitive framework. Therefore, it is dangerous to apply such solutions in domains where the users must adhere to the Law of Armed Conflict and consider other ethical constraints. Finally, a problem with ML-approaches is that the ML-system will learn independently on each dataset and potentially create one structure from one dataset that is incompatible with the next result. As a result, the approach may lead to the development of additional silos of knowledge.

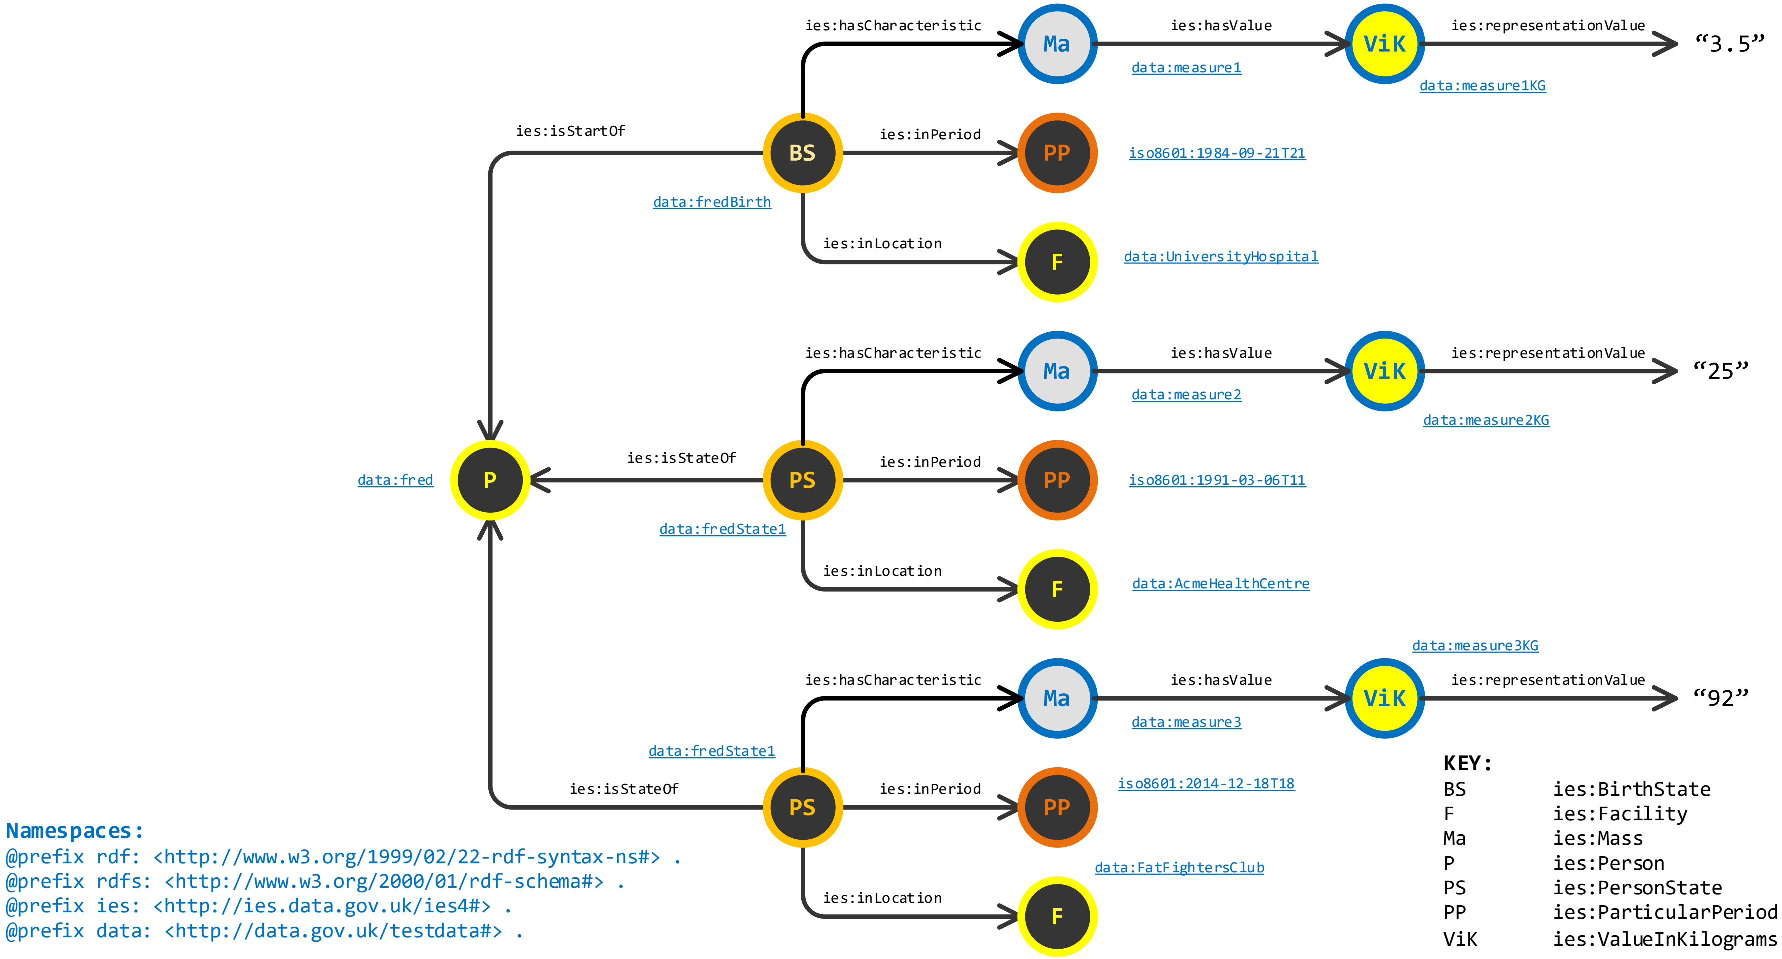

Figure 3-2: Example IES Schema.

3.4.1.2 With New Technologies, Doesn't That Just Lead to New Complexities?

Question: New approaches and technologies often needed to develop ontologies seem complicated – is it possible to achieve robust solutions that people can engage with?

Reply: Ontologies are based around common language using the subject-predicate-object construct which is infinitely less complex than abstract computer coding and allows for input from Subject Matter Experts (SMEs) in a simpler and more meaningful way.

Question: Does this mean I have to move to new data storage solutions and complexity?

Reply: You do not have to move to new data storage solutions. In fact, your data storage options are expanded. Ontologies allow for a data centricity vs. application centricity which means data is not 'held hostage' by an application but rather can be utilised through different applications

Question: By constraining my database solution to a graph database, will I incur latency problems for complex queries and updating?

Reply: The best measure seems to be to use a modular approach to deal with complexity. A key skill in modelling is knowing when enough is enough and focusing on specific needs (don't boil the ocean) – use decision-support questions to scope to challenge. For smaller databases that do not require frequent updating, a graph database is likely the optimal solution. However, for implementations such as C2 of mobile military entities, they are less practical. Fortunately, there are several open source and commercially available tools and data base enterprises to allow integration of the ontology with federated and mixed database formats. One way to deal with potential complexity and query latency is to constrain the knowledge model to an OWL2 profile such as a Query Language (QL) which allows for efficient linkage of the model with relational databases.

3.4.1.3 Are KR Methods Robust?

Question: Can KRR respond to new requirements?

Reply: Use the modular approach, combining a core with related reference and application ontologies. The core remains fairly rigid in order to ensure interoperability between users of related reference and application ontologies. The more specialised modules are expandable or adjustable as the mission owner requires them to be. For governance, establish a responsive governance structure. It should entail both coordination and development processes to improve the collection of modules as well as the technology in place to conduct automated validation of ontology alignment and logic as modules evolve.

Question: I hear that ontologies constrain our vernacular because we are forced to use terms in the ontology that are not common to our domain. Is that true?

Reply: Formal ontology languages provide tools to overcome this problem and allow many different terms for a similar entity to co-exist in the model. There is some need to adhere to the overall standard, but the benefits gained in interoperability should far outweigh the cost.

3.4.1.4 Are Specialist Skills and Expertise Required?

Question: What skills, expertise and experience are required in teams to develop and deploy KRR methods?

Reply: KRR approaches, utilising ontologies and/or other structures to provide symbolic representations represent a different approach to traditional schema based methods required by relational databases. This report has outlined the skills and understanding required to effectively use knowledge reasoning and representation methods and has given examples where they are in use.

The discipline of knowledge engineering, well understood to support the implementation of Symbolic AI / expert systems, may now require strengthening across NATO nation capabilities. This may require the appointment of knowledge engineers in organisations. Future knowledge engineers would assume responsibility for defining core and domain specific ontologies, establishing systems that SMEs can maintain, providing the conceptual leadership for future knowledge-based systems and helping to navigate the route to integrating with existing knowledge systems. Knowledge engineers require knowledge and experience of KRR methods and technologies highlighted in this report, but they also require the skills to interact with SMEs with unstructured and structured techniques to allow them to dissect and represent a particular domain in the most appropriate manner.

3.4.2 The Strengths and Weaknesses of Knowledge Representation and Reasoning Approaches

The following list acknowledges some of the advantages of symbolic approaches:

- Given the right tools knowledge representations and rule systems can be coded and maintained by knowledge engineers who are operational SMEs themselves or who can naturally form a team with an SME. Problems frequently arising from miscommunications between software engineers and operational experts can be avoided.
- While ML-techniques rely on great amounts of data, symbolic approaches do not rely on data at all but only on the knowledge of subject matter experts. Especially in the military domain, we often face the problem that sufficient data for training ML-models do not exist or are not accessible. This is never a problem for a symbolic approach.
- While the update of an ML-model often requires a complete new training of the model – with respectively high costs in time and computational power – the update of a symbolic model is rather straightforward.
- Inferences that are drawn from a knowledge base are, in principle, comprehensible. They are justified by the applied rules of (rational) reasoning. Therefore, the output of a symbolic reasoning system can, in principle, be explained to the operator.
- Knowledge representations can be Semantic References Models (SRMs) for information that is to be exchanged among partners. They can enable semantic interoperability.

In addition to the positive points outlined here, some challenges must also be considered:

- The size of a knowledge representation provides more chances for the introduction of inconsistencies. As inconsistencies are necessarily false and from something false literally everything can be concluded, they pose a serious problem. Either the consistency of a knowledge base has to be guaranteed or an approach to dealing with inconsistencies has to be found (which, however, almost certainly will affect inferencing in general).
- A good way to overcome the challenge is to follow a modular approach and to extend core systems and ontologies for well-defined, bounded contexts – those that follow the rules of domain-driven design [13].
- Reasoning with large knowledge bases is computationally demanding and time consuming. Worst case complexity is high. This issue is usually resolved by constraining the complexity of the knowledge model. For example, by using an OWL2 QL profile, some properties can be constrained or disallowed such that the model can be efficiently linked to relational database and query response time can be greatly reduced.

- The effort for manually coding and maintaining knowledge bases can be quite high.
- This challenge is often addressed by integrating AI techniques, such as NLP calibrated with a basic ontology to generate recommendations for term expansion. Another measure could be to apply ML for the automatic learning of ontologies/knowledge graphs [14].
- Facts and rules must be known before they can be coded. Symbolic approaches are not suitable for exploration. (In comparison, ML-techniques or other statistically based approaches can be suitable for exploration.)
- The establishment of a knowledge representation as a standard reference for interoperability requires collaboration between the relevant stakeholders and, thus, further effort in its specification. This challenge is addressed by ensuring the development team contains both knowledge engineers and domain SMEs. The development must also be iterative to start small and build the knowledge base over time.
- Automatic reasoning depends on coded facts and rules. In order to infer all potentially relevant conclusions from a knowledge base, it must be complete, that is, contain *all* potentially relevant facts and rules. Completeness is hard to guarantee.
- Data can be vague or even false. A reasoner has to be able to deal both with fuzzy and potentially false data if sensor data are to be considered.

Thus, while symbolic methods offer the chance to encode SME knowledge and offer comprehensible inferences, they don't suit every situation and could incur an overhead in maintaining knowledge bases and rulesets.

3.5 SUMMARY

In this chapter, we discussed semantic standards and tools for enabling information integration and arriving at a common understanding of heterogeneous information. To these standards belong the open standards of the W3C Semantic Stack, among others. We described example of ontologies in use or development and provide four paradigmatic examples for the application of knowledge representation techniques, namely:

- The DICO as an ontology for supporting reasoning in defence intelligence;
- The WISDOM platform for studies in information integration, sense-making as well as specifying and exploiting expert and domain knowledge; and
- An experimental system for information filtering, aggregation and enrichment within the INSANE environment.
- An ontology driven standard for Information Exchange – the UK Intelligence Community (UKIC) Information Exchange Standard (IES).

Finally, we discussed common concerns regarding knowledge technologies and the strength and weaknesses of the knowledge representation approach.

Key takeaways from this section include:

- Ontologies and information models can serve as semantic reference models for the integration of distributed information as well as semantic interoperability solutions.
- Some information exchange models, like the NIEM, are rather definitions and message formats and cannot count as ontologies. Not every model that has been designed to enable information exchange can serve as a proper semantic reference model.

- The W3C Semantic Web Stack is a set of open standards for defining and exchanging knowledge. Military developments can profit from referring to these standards and standards compliant open software.
- Contrary to ML systems, KBS do not demand large amount of training data and are therefore the choice, when data is not sufficiently available.
- Contrary to ML systems, KBS can ensure explainability of their output.
- Scalability can be an issue for KBS. It is therefore recommendable to follow a modular and domain-driven design approach, with manageable solutions for bounded contexts. This facilitates also adaptation to new requirements.

ML-techniques and knowledge representation techniques can be combined, e.g., to automatically create knowledge representations (ontology learning) or to ensure explainability of ML systems, e.g., Explainable AI (XAI).

3.6 REFERENCES

[1] Dechand, M., Gautreau, B., Sikorski, L., Trautwein, I., Bouvier, E. and Khimeche, L. (2019). Development of an Air Operation eXtension with the (Future) C2SIM Standard. 2019 NATO Modelling and Simulation Group Symposium.

[2] Pullen, M., Corner, D., Singapogu, S., Blais, C. and Reece, D. (2019). Command and Control System to Simulation System Interoperation: Development of the C2SIM Standard. Naval Postgraduate School.

[3] Blais, C., Galvin, K. and Hieb, M. (2004). Coalition-Battle Management Language (C-BML) Study Group final Report, Simulation Interoperability Standards Organization.

[4] NIEM. (2020). National Information Exchange Model (NIEM). Retrieved 17/11/2020, 2020, from https://www.niem.gov.

[5] Arp, R. Smith, B. and Spear, A.D. (2015). Building Ontologies with Basic Formal Ontology, MIT Press.

[6] W3C. (2020). Semantic Web Standards. Retrieved 17/11/2020, 2020, from https://www.w3.org/2001/sw/wiki/Main_Page.

[7] Glimm, B., Horrocks, I., Motik, B., Stoilos, G. and Wang, Z. (2014). HermiT: An OWL 2 Reasoner. Journal of Automated Reasoning 53, pp. 245-269.

[8] Sirin, E., Parsia, B., Grau, C., Kalyanpur, A. and Katz, Y. (2007). Pellet: A Practical OWL-DL Reasoner. Journal of Web Semantics 5(2), pp. 51-53.

[9] Roy, J. (2020). WISDOM – A Platform for Proof-of-Concept R&D in Data and Information Integration, Fusion and Analytics. Defence Research and Development Canada (DRDC).

[10] Bau, N., Endres, S., Gerz, M. and Gokgoz, F. (2018). A Cloud-Based Architecture for an Interoperable, Resilient, and Scalable C2 Information System, IEEE, pp. 1-7.

[11] Schmitz, H.-C., Bau, N., Endres, S., Gerz, M., Käthner, S. and Mück, D. (2018). A Newsfeed for C2 Situational Awareness. International Command and Control Research and Technology Symposium (ICCRTS) Pensacola/Fla.

[12] Ahmed, O., Gökgöz, F. and Schmitz, H.-C. (2020). Unsupervised Creation of Profiles for Information Management. International Command and Control Research and Technology Symposium (ICCRTS) Southampton, UK.

[13] Evans, E.J. (2003). Domain-Driven Design: Tackling Complexity in the Heart of Software, Addison-Wesley Professional.

[14] Hamilton, W. (2020). Graph Representation Learning, Morgan & Claypool.

Chapter 4 – CURRENT RESEARCH THEMES FOR KNOWLEDGE REPRESENTATION AND REASONING

In the section to follow, we examine a selection of active research themes for KRR. While relevant to a variety of different computational systems, these themes reflect the points at which KRR overlap with AI research, including the development of autonomous agents that can serve as teammates to humans by accomplishing different forms of recognition, interpretation, and reasoning over sensory inputs from the world around them, including language. We begin with an overview of text and joint text and image/video analytics (Section 4.1), turn to interfacing with humans through natural language (Section 4.2), then discuss the representation of causality (Section 4.3), followed by explainability and trust in inferencing (Section 4.4). We close with a summary and a presentation of the outlook, as well as recommendations for developing skills and expertise for tackling these ongoing challenges and dynamic technology landscape.

4.1 MULTI-MODAL KNOWLEDGE REPRESENTATION – DEALING WITH TEXT, IMAGES, AND BEYOND

Text and image/video data, among other unstructured data sources, share common characteristics that make exploiting the data difficult, especially in an army tactical environment where there is likely to be noisy, conflicting data. The sheer volume of text and image/video data represents one challenge: given the “fire hose” of data streaming in, how do we constrain what undergoes further analysis? A second challenge is that text and image/video analytics have advanced as two separate fields, which are only beginning to come together to establish shared structures for exploiting text and video jointly. Nonetheless, both fields have normally sought to extract relevant features from the input and classify the data accordingly in order to impose any kind of “meaning” on unstructured text and video data.

There are three broad types of classification approaches, which often differ in the amount of knowledge injected into the system: supervised, unsupervised, and semi-supervised. Supervised approaches tend to take a more knowledge-based approach to classification, often relying on large amounts of manually annotated data in combination with a taxonomy or ontology of words and concepts. Unsupervised approaches tend to take a more statistical approach to classification, relying on the assumption that “you can know a word by the company it keeps,” and similarly for images, that you can know a pixel (or sub-space features) by examining distributions and co-occurrence patterns. This section examines the trends in text analytics, including vector space representations, which have been adopted in text and image analysis, but also in joint modelling of data from a variety of disparate sensors.

4.1.1 Symbolic Approaches to Text Analytics

Largely drawing upon research in AI and Linguistics, one approach to these problems uses symbolic, structured methods that often rely on the existence of manual annotations. The manual annotations consist of tags organising what are thought to be the relevant features of the data, and these may also be organised in a taxonomy or ontology of words, concepts, or linguistic meta-categories, such as syntactic categories. At this point, technologies have advanced such that many syntactic patterns can be automatically detected for high-resource languages at relatively high accuracies. As a result, morpho-syntactic analysers are quite dependable for languages like English [1]. Furthermore, based on syntactic notions like subject and object, which can be fairly reliably captured (again, for higher-resource languages) in phrase structure and dependency parses [2], some semantic correspondences can be exploited in languages where certain semantic roles are often associated with certain syntactic slots. For example, in English, the syntactic subject is often an actor, while the syntactic object is often an undergoer: *David*$_{\text{ACTOR}}$ *kicked the ball*$_{\text{UNDERGOER}}$. Exploiting such patterns between syntactic realizations and semantic features is of primary interest to automatic Semantic Role Labelling (SRL), described in greater detail in the paragraph to follow [3].

Unfortunately, while nouns and verbs can be identified, many concepts must be inferred, such as causation and intention. One way to encode such concepts is within a lexical resource, such as a taxonomy or ontology of words/concepts. When associated with certain tags used in an SRL annotation schema, patterns can emerge, such as the fact that communication events generally have two co-agents as participants: *Susan*$_{\mathrm{AGENT}}$ *chatted with Rachel*$_{\mathrm{CO\text{-}AGENT}}$. A variety of linguistic resources have been developed to capture generalizations about what types of participants are involved in a given type of event, and each has been used to annotate data and train SRL systems: The Proposition Bank [4], ERE [5], VerbNet [6], and FrameNet [7]. The SRL can, in turn, be used in more complex applications like information extraction and summarization, and has even been shown to improve machine translation [8].

Although each of these resources was created with a slightly different goal in mind, they are surprisingly compatible in that each offers a unique strength, given the varying level of semantic specificity represented in each. Thus, ideally these resources could be used together or interchangeably, depending upon what type of information was desired, and the annotated data created for each could be combined into a larger, more diverse training corpus. See Section 2.3.3 for a description of the Rich Event Ontology, which unifies SRL resources.

4.1.2 Vector Space Models of Text

In contrast to knowledge-based approaches, many statistical approaches can be applied in an unsupervised fashion, which requires less overhead in the form of creating annotation schemas and performing the annotation. The fundamental problem is to learn meaning and usage of words in a data-driven fashion from a corpus without prior linguistic knowledge. Statistical approaches to this problem, perhaps not surprisingly, are primarily based on word frequencies and distributions in large corpora. Such statistical measures can allow insight into meaning, for, as Firth [9] put it, "You shall know a word by the company it keeps". Firth's quote captures the notion that word meaning is necessarily context-dependent, and that words with similar meanings will share similar contexts. This notion has given rise to a variety of influential vector space models of word and document meaning. A vector space model is an algebraic model used for representing text (but in theory, it can be used to represent any object) as vectors of identifiers [10]. In a vector space model, each dimension corresponds to a separate term, and if a term occurs in the document, its value in the vector is non-zero. The precise way in which the value, known as a term weight, is computed depends upon the specific model.

Another influential family of statistical learning models has also been introduced and taken up rapidly in NLP: neural networks [11]. Artificial neural network learning is inspired by biological neural networks in the central nervous systems of humans and animals. Although there is no single formal definition of an artificial neural network, it generally can be thought of as a network of interconnected nodes or "neurons" that can exchange "messages" between each other. The connections between nodes have numeric, adaptive weights that can be tuned based on experience, much in the way that there are different connections strengths between neurons in the nervous system, and frequently activated paths have stronger connections. There is firstly a set of input neurons, which may be activated by the words in a document (or the pixels in an image). Depending upon the weights or strength of connections between the input nodes and connected nodes, the activation will be passed on to other neurons. This process is repeated until an output neuron is activated. Like term weights, the weighting of the connections depends upon the specific model examined. A popular neural network model is Word2Vec [12]. Word2Vec computes vector representations for words that are learned by neural networks, which arguably preserve some linear regularities among words and is less computationally expensive on large data sets than Latent Dirichlet Allocation (LDA).

An important assumption of Word2Vec is that "similar" words tend to be close in the vector space, but of course words can have multiple types and degrees of similarity. Somewhat surprisingly, one can use "word offset techniques," where simple algebraic operations are performed on the word vectors to find analogical relations. For example, "what is the word that is similar to *woman* in the same sense as *king* is to *man*?" To answer the question a:b, c:d where d is unknown, one can find embedding vectors a, b, c and compute d = b – a + c. In other words, King – Man + Woman = Queen. Figure 4-1 depicts this and a syntactic relation (singular/plural) in an idealised, two dimensional vector space representation.

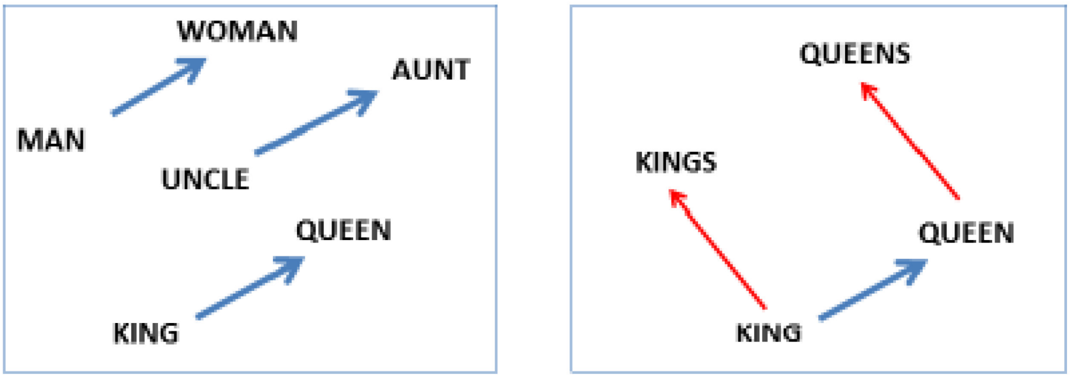

Figure 4-1: Semantic and Syntactic Vector Offset Relations [12].

4.1.3 Combined Vector Space and Knowledge-Based Approaches to Text Analytics

Although promising, vector space representations generally have been criticised for being ad hoc and prone to the production of low-quality vectors since training text often contains incomplete, ambiguous information or simply not enough data. Essentially, although vector space models are able to capture topical similarities between words found in a text, what type of similarity is captured can vary wildly from word to word and text to text. Accordingly, the efficacy of using vector offset relations may depend upon what type of relation is sought. Vylomova, Rimell et al. [13] explore a range of semantic and syntactic relations and find that, in general, morphosyntactic relations are more accurately represented than lexical semantic relations. This has led researchers to find approaches that use vector space representations, but also leverage some type of knowledge to produce word representations of higher quality. Several notable combined approaches are discussed in the paragraphs to follow.

Xu, Bai et al. [14] explored an early approach of this type, incorporating information from a knowledge graph containing information of two types: relational knowledge and categorical knowledge. Relational knowledge encodes relationships between entities so as to differentiate word pairs with analogical relationships (e.g., is-a, part-of, child-of relations). Categorical knowledge encodes attributes and properties of entities, according to which similar words can be grouped into meaningful categories (e.g., gender, location). The authors test a system augmented only with categorical knowledge (C-Net model), augmented only with relational knowledge (R-Net model), and one augmented with both (RC-Net model) against the Skipgram model. The systems are evaluated on the same semantic/syntactic analogies tests from Mikolov, Yih et al. [15]. The authors find that while both knowledge-powered models outperform Skipgram, the full RC-Net model yields the largest improvements. Interestingly, the authors find that while incorporating relational knowledge can give rise to higher accuracy on all types of analogical reasoning tasks, incorporating only categorical knowledge can cause increased performance on semantic analogies, but decreased performance on syntactic analogies. Intuitively, this makes sense given that syntactic analogies represent only one type of relational knowledge, and other semantic relations between syntactic elements will only cause confusion when trying to pinpoint syntactic relations.

In Faruqui, Dodge et al. [16], the authors explore retrofitting existing word vectors to semantic lexicons, including the Paraphrase Database (PPDB) [17], WordNet [18] and FrameNet [7]. The authors explore refining vector space representations using relational information from each of these semantic lexicons – essentially encouraging words linked in the lexicon to have similar vector representations. This idea is captured graphically in Figure 4-2.

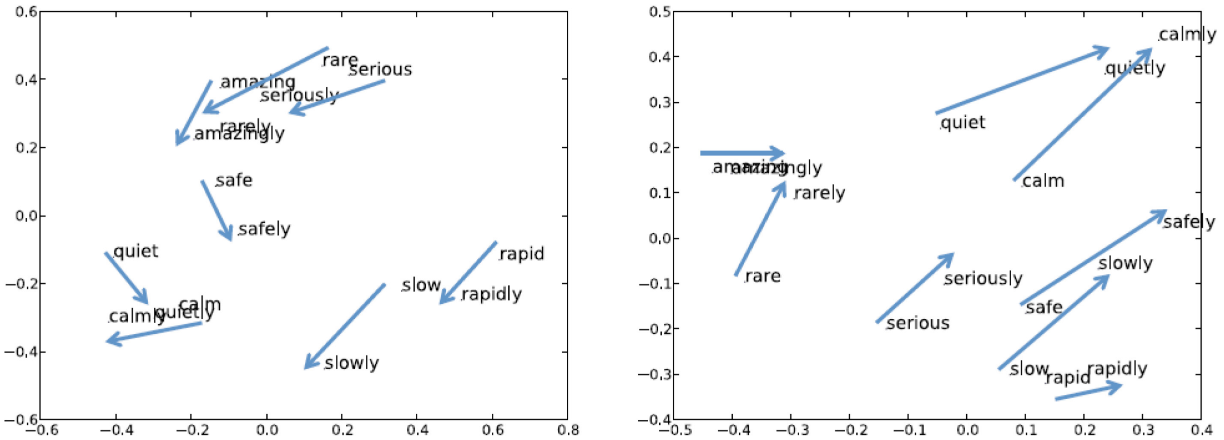

Figure 4-2: Two Dimensional Projections of 100 Dimensional Vector Pairs Holding the "Adjective to Adverb" Relation, Before (Left) and After (Right) Retrofitting [16].

Notably, the authors offer this as a post-processing step; thus, they are agnostic as to what type of vector representation is used. The authors test retrofitting a variety of different vector models with each of the semantic lexicons, evaluating each model on word similarity, syntactic relation, synonym selection and sentiment analysis tasks (note the variety of applications to which vector space representations have been applied). The authors find that retrofitting with any of the lexicons (except FrameNet) offers improvements on the word similarity/synonym tasks as well as the sentiment analysis task. The poor performance of FrameNet is explained by the fact that FrameNet doesn't group words according to semantic or syntactic similarity, and instead according to shared real world domains – a very abstract type of similarity. Similarly, the authors find no improvement on the syntactic relation task, given that the type of knowledge encoded by the selected lexicons is semantic only. This is consistent with what Xu, Bai et al. [14] noted with respect to the poor performance on syntactic relation tasks by the model augmented only with categorical knowledge. It is clear that one must carefully select the type of knowledge injected into the system, and weigh this with respect to the planned task.

While the existing models of vector space word representations are based solely on linear contexts, Levy and Goldberg [19] explore the use of more structured contexts by making use of dependency parses. Dependency parses provide syntactic dependency relations between words in a sentence. For example, an adjective depends upon a noun that it is modifying. An example dependency parse is shown in Figure 4-3. Note that in this parse, there is prepositional attachment ambiguity: from the perspective of a computer, "with telescope" could either be the instrument of discovery (modifying "discover") or something the star has (modifying "star"). The dependency parse clarifies the appropriate interpretation.

In this research, the authors derived the context for the vector space representation by considering the type of dependency relation between the head (thing depended upon) and modifier (dependent). This allows the vectors to captures relationships between words that are far apart (such as *discover* and *telescope*) and therefore out of reach in a linear context model. It also allows for the filtering out of coincidental contexts which are within the window, but are not directly related (such as *star* and *telescope*). Perhaps most importantly, this method provides context information that is typed: indicating, for example that *scientists* are subjects of *discover* and *stars* are objects. In a qualitative analysis of target words and their 5 most similar words, comparing the dependency-based model and Continuous Bag Of Words (CBOW) models with windows of both 2 (CBOW2) and 5 words (CBOW5), one can quickly see the difference between what is "similar" to the different models. For example, using the target word *Florida*, the most similar words using the CBOW model are largely cities within Florida or the abbreviation *Fla*, whereas the dependency-based model lists only other U.S. states. The authors also compare their model to CBOW models in a quantitative fashion, testing each model's ability to rank "similar" pairs of words (defined as functional similarity) over "related" pairs of words (defined as topical similarity), given a set of human-annotated "related" pairs or "similar" pairs. The dependency-based model achieves the greatest success in this task, over both the CBOW2 and CBOW5 models. Thus, this model represents a promising approach to improving vector space representations with syntactic information.

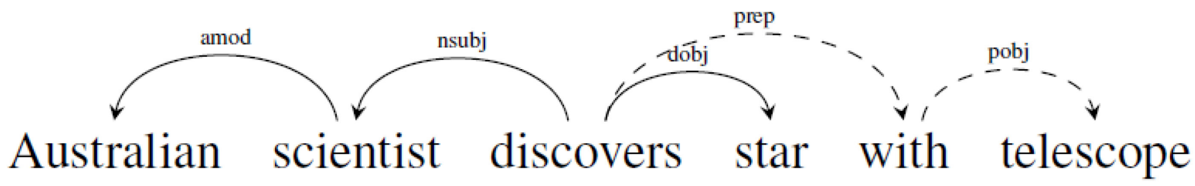

Figure 4-3: A Dependency Parse, Illustrating Syntactic Dependencies between Words [19].

4.1.4 Joint Modelling of Text and Imagery

Computer vision research has also found value in using neural network learning of vector space models, and there is naturally a growing body of research attempting to model an image and an accompanying, explanatory text (such as a caption or object labels) "jointly." Presumably, any model that integrates features from both text and images is a multi-modal, "joint" model. The extent and the point at which the model truly integrates text and image representations vary, as exemplified in the research described below.

Current research [20] has aimed at providing fuller descriptions of an entire scene by jointly modelling entire images and descriptive phrases, as opposed to objects within images and object labels. A notable approach [21] uses a truly multi-modal vector space by embedding fragments of images (containing objects) and fragments of sentences (dependency tree relations) into a common vector space. This research builds upon previous work [22] that also used dependency tree relations in a multi-modal vector space representation of images and captions, but is novel in that it uses fragments of sentences and images. The goal of this system is to retrieve relevant images given a sentence query, and conversely, relevant sentences given an image query. Potential image-sentence pairs are scored based on how confidently certain sentence fragments can be matched to some fragment in the image. Like Levy and Goldberg [19], the authors choose to use dependency relations in their sentence fragment embeddings. Their image fragment embeddings are based on a neural network model that is pre-trained on ImageNet [23]. The system is evaluated on image-sentence retrieval performance using datasets of 1,000, 8,000, and 30,000 images respectively, wherein each image is annotated using Amazon Mechanical Turk with 5 distinct sentence descriptions. Their model outperforms previous methods applied to the same datasets, which the authors take as evidence of the fact that fragment object representations are more informative than global scene representations. The authors also conclude that dependency tree relations outperform bag of words and bigram representations since dependency relations provide useful structure that the neural network takes advantage of. There is certainly room for improvement in this system, given certain limitations. First, a single phrase describing a single visual entity can be split across multiple sentence fragments (e.g., "black and white dog" has a dependency relation between "black" and "white," and a distinct dependency relation between "white" and "dog," effectively splitting this modifier across distinct fragments). Secondly, this model does not take into account single phrases describing multiple visual entries (e.g., "Three children playing"). Finally, the object detection system lacks spatial information; therefore, it may mistakenly detect, for example, multiple people inside one person.

Beyond just text and imagery, joint modelling research has gained traction over the course of the DARPA program "Active Interpretation of Disparate Alternatives" (AIDA). The goal of this program is to take multiple disparate unstructured data sources, such as text and speech of multiple languages, but also images, video and other sensor data, and convert these sources into a single, common semantic representation from which knowledge and hypotheses with confidence measures can be derived. In order to undertake this challenge, performers in this program have utilised a multi-modal embedding space, where various data sources are modelled in a single embedding space, or in a hypergraph [24]. There is promise for these approaches to be combined with symbolic, knowledge-based approaches, like the combined approaches to word embeddings described previously.

4.2 CONSIDERATIONS FOR HUMAN INTERFACING – NATURAL LANGUAGE INTERFACES

One primary purpose of KRR capabilities is to enable agents to both understand and potentially communicate about the world around them as another human might. As humans communicate in language, there are significant bodies of research in dialogue systems, computational semantic representations of natural language, and research supporting the recognition of speech acts, or what someone is attempting to do with a particular utterance beyond its basic content. All of these areas of research support natural language interfaces with agents. The following provides a brief overview of research in these areas.

4.2.1 Dialogue Systems

Task-oriented spoken dialogue systems – the goal of which is broadly to identify a user's intents and then act upon them to satisfy that intent – have been an active area of research since the early 1990s. Broadly, the architecture of such systems includes:

- Automatic Speech Recognition (ASR) to recognise an utterance;
- A Natural Language Understanding (NLU) component to identify the user's intent; and
- A dialogue manager to interact with the user and achieve the intended task [25].

The meaning representation within such systems has, in the past, been predefined frames for particular subtasks (e.g., flight inquiry), with slots to be filled (e.g., destination city) [26]. In such approaches, the meaning representation was crafted for a specific application, making generalizability to new domains difficult if not impossible. Current approaches still model NLU as a combination of intent and dialogue act (e.g., a question or statement) classification and slot tagging (identifying semantic entities of interest in an utterance), but many have begun to incorporate Recurrent Neural Networks (RNNs) and some multi-task learning for both NLU and dialogue state tracking [27], [28], the latter of which allows the system to take advantage of information from the discourse context to achieve improved NLU. Substantial challenges to these systems include working in domains with intents that have a large number of possible values for each slot and accommodation of out- of-vocabulary slot values (i.e., operating in a domain with a great deal of linguistic variability).

Thus, a primary challenge today and in the past is representing the meaning of an utterance in a form that can exploit the constraints of a particular domain but also remain portable across domains and robust despite linguistic variability.

Although human-robot dialogue systems often leverage a similar architecture to that of the spoken dialogue systems described above, human-robot dialogue introduces the challenge of physically situated dialogue and the necessity for symbol and action grounding, which generally incorporate computer vision. Few systems are tackling all of these challenges at this point (but see Ref. [29]).

4.2.2 Semantic Representation of Natural Language

There is a long-standing tradition of research in semantic representation within NLP, AI, as well as theoretical linguistics and philosophy (see Ref. [30] for an overview). In this body of research, there are a variety of options that could be used within dialogue systems for NLU. However, for many of these representations, there are no existing automatic parsers, limiting their feasibility for larger-scale implementation. Two notable exceptions with a body of research on automatic parsing are Combinatory Categorical Grammar (CCG) [31] and Abstract Meaning Representation (AMR) [32]. CCG parsers have already been incorporated in some current dialogue systems [33]. Although promising, CCG parses closely mirror the input language, so systems making use of CCG parses still face the challenge of a great deal of linguistic variability that can be associated with a single intent. In contrast, AMR abstracts from surface variation; thus, AMR may offer more regular, consistent parses in comparison to CCG. AMR is being investigated for use in dialogue systems on-board robots used for search and navigation tasks [34], [35].

4.2.3 Speech Acts and Dialogue

In order to engage in dialogue, an interlocutor must interpret the meaning of a speaker's utterance on at least two levels, as first suggested by Austin [36]:

1) Its propositional content; and
2) Its illocutionary force.

While semantic representations have traditionally sought to represent propositional content, speech act theory has sought to delineate and explicate the relationship between an utterance and its effects on the mental and interactional states of the conversational participants. Speech acts have been used as part of the meaning representation of task-oriented dialogue systems since the 1970s [37], [38], [39]. For a summary of some of the earlier work in this area, see Traum [40]. Although the refinement and extension by Austin of Searle's hypothesized speech acts remains a canonical work on this topic [36], [41], there have since been a number of widely used speech act taxonomies that differ from or augment this work, including an ISO standard [42]. Nevertheless, these taxonomies often have to be fine-tuned to the domain of interest to be fully useful.

There is a growing interest in representing various levels of interpretation in existing meaning representation frameworks, and in AMR in particular. Bonial, Donatelli et al. [34] present Dialogue-AMR, which augments standard AMR, representing the content of an utterance, with speech acts representing illocutionary force; the authors present the fully annotated corpus of human-robot dialogue [43], Dial-AMR, with parallel standard AMR and Dialogue-AMR mark-up. Bastianelli, Giuseppe et al. [44] present their Human Robot Interaction Corpus (HuRIC) following the format of AMR. This corpus is comprised of paired audio interactions and transcriptions. Though all text is annotated in the format of AMR, AMR is significantly altered by incorporating detailed spatial relations, frame semantics [45], and morphosyntactic information. Shen [46] further presents a small corpus of manually annotated AMRs for spoken language to help the parsing task. The study presents findings that while AMR offers a clean framework for the concepts and relations used in spoken language, the mapping between AMR and computer-interpretable commands is not trivial, especially in the case that very little of training data is provided.

Such work is paralleled by a more sustained recognition of and interest in the multifunctionality of utterances in dialogue across the dialogue literature [47], [48], [49]. O'Gorman, Regan et al. [50] present a Multi-Sentence AMR corpus (MS-AMR) designed to capture co-reference, implicit roles, and bridging relations. Though not strictly speech acts, the interconnected approach to meaning that this corpus annotates is directly relevant for deducing illocutionary force in a dialogue context. Kim, Kane et al. [51] similarly describe an annotation schema designed to capture discourse inferences via underlying semantic scope relations. Hajicova [52] outlines an argument for modelling information and discourse relations explicitly in meaning representations. Though none of these proposals looks at illocutionary force directly, the recognition that meaning representations for dialogue need to be expanded to capture levels of interpretation beyond the propositional content is growing in NLP.

4.3 CAUSALITY AND CAUSAL MODELS

Causality is an intuitive concept commonly used for understanding and explaining processes and courses of events, for predicting future events, for intervening into courses of events and for judging on actions and the responsibility of actors. We ask for causes, e.g., when we investigate *why* a building has collapsed, what *causes* lung cancer or *makes it more probable*, *how* we can prevent a disease from spreading, and whether someone can be held *responsible* for an accident. Causal thinking is related to conditional thinking – *what* will happen, *if* we do this? – and counterfactual thinking – if the cause had not occurred, the effect would not have occurred [53]. It is, thus, an important means of hypothetical reasoning.

Although intuitively clear, the concept of causality is hard to define and more than once it has been eliminated from scientific discourse [54]. In statistics, the asymmetric concept of causality has been largely given up in favour of the symmetric concept of correlation. The elimination of causality, however, has let to problems regarding the interpretation of observational data that has not been generated by a randomised controlled experiment. An example is Simpson’s paradox which occurs when an association between a pair of variables is consistently converted in each subset of a partitioned population. Structural causal models can contribute to explaining such anomalies and resolve the question on whether a decision should be based on the statistics of the entire population or its partitions [55].

A Structural Causal Model (SCM) consists of two sets of variables **U** and **V** and a set of equations that assigns each member of **V** a value based on the other variables in the model. We call the variables in **V** the endogenous variables of the model. Their values are determined by the other variables via the equations. We call the variables in **U** the exogenous variables of the model. Their values are not depended on other variables. Instead, they are determined by a given probability distribution. Given that the SCM can be represented by a Directed Acyclic Graph (DAG), the probability distribution over **U** determines the values of all variables in the model. A variable X is a direct cause of a variable Y if and only if X appears in the equation that assigns Y’s value (thus, only variables in **V** have causes.) [56].

A DAG together with a probability distribution can be considered a causal Bayesian net. An SCM as defined above is a probabilistic causal model. It is also possible to define SCMs as deterministic models [57]. Causal models can be manually coded or, in principle, be learned from data (“causal discovery”) [58]. Causal models should be both transparent and testable, that is, they should be empirically falsifiable. The usage of graphical models contributes to transparency and compactness of causal models.

A great influence for the reintroduction of causality into scientific discourse and the application of SCMs for reasoning is the work of Pearl [59], among the works of other authors. Pearl combines graphical models with his *do*-calculus and thereby allows for the identification and de-confounding of confounding variables, the modelling of interventions, and, finally, the evaluation of conditional and counterfactual statements. The reasoning process implemented in Pearl’s framework is depicted in Figure 4-4 which is taken from Ref. [58]. Causal assumptions are coded by a graphical model. Based on that model an estimand for a query can be computed that “provides a recipe for answering the Query from any hypothetical data”. Given the estimate, the query can be answered (along with statistical estimates) by referring to actual data. Finally, fit indices can be computed to evaluate the measure the compatibility of the assumption and the data and, thus, evaluate the causal model.

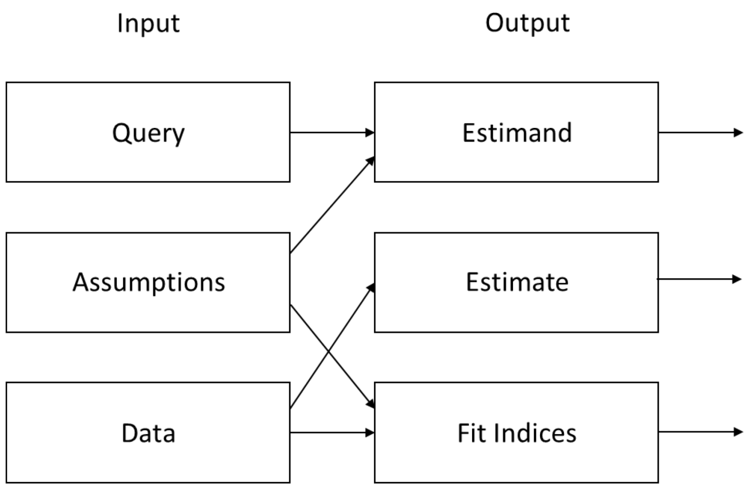

Figure 4-4: SCM Inference Engine According to Pearl [58].

Causal models and the *do*-calculus have been proven very useful i) for modelling interventions and hypothetical reasoning (including counterfactual reasoning) and ii) for merging and integrating different data sets and transporting solutions to different contexts, thereby contributing to their external validity. We assume that there is a great potential for applying causal models in the military domain, in particular for improving robust decision support.

4.3.1 Causality in Natural Language Processing

Causation exists as a psychological tool for understanding the world independently of language, but language can be used to express causation both explicitly and implicitly. To reason over language and the world around them, autonomous systems and agents must be equipped with some causal knowledge. What is more, vast amounts of language data available on the web and even in the form of film scripts [60] can be used to teach systems common-sense causation, in what has been called causative discovery. This information includes causation that is counterfactual, but also more nuanced conditions that must be in place for an event to happen, sometimes called the "causal complex" [61]. For example, to bake bread, the ingredients must of course be combined and baked, otherwise the baking event would not occur, but one must also presumably shop for the ingredients, have a heat source and appropriate baking vessel, etc. This brings to light some of the main challenges of causal discovery – what factors should be enumerated as relevant, and how do we know what these factors are, given that often times many factors are so ingrained in everyday rituals that they are rarely mentioned explicitly.

To facilitate recognition of various types of causation, several different types of resources have been established that draw upon NLP techniques in different ways. First, there are lexical resources that spell out the words that encode some kind of causation; these include words like cause and make, but also words which "lexicalize" causation, like murder, meaning cause to die [7] [62]. Secondly, there are resources that provide manual annotations over text corpora that indicate the causal relationships between contiguous sentences [63]. In these cases, the causal relations may be made explicit by discourse connectives, such as therefore, or they may be entirely implicit: It was raining. I put on my coat. Other types of manual annotations focus on both causal and temporal relations between two events expressed in text [64], or between two contiguous sentences in narrative text [65]. All of these manually crafted resources can be used as ML training data, such that the causative information can be predicted and understood in previously unseen text. However, pinpointing causation remains difficult for even human annotators, and many of these projects are plagued by low or moderate agreement between human raters; thus, unsurprisingly, there is also room for improvement in system performance.

Once trained on the manually curated data, these systems can and have been used to support causal discovery in previously unseen text. However, as these systems are finding causation in text, there is a bias towards learning about the events that are mentioned in text, and not the everyday, common-sense causation that is implicit. This is the inevitable reporting bias in language, as people tend to discuss and write about the atypical over the expected [66]. For example, if we want a system to learn what objects are heavy and light from looking at text corpora, it may come to the conclusion that trucks are light because “light trucks” are mentioned frequently, whereas the default assumption is that a truck is heavy. The recently developed ATOMIC resource attempts to overcome this bias by crowd-sourcing common-sense causation around an initial set of 24,000 common events extracted from text [67]. For example, participants are queried about what would cause Person X to hug Person Y, and what effect the hugging event would have on both Person X and Y. With the resulting knowledge graph of over 300,000 nodes, it has been shown that neural models can acquire simple common-sense capabilities and reason about previously unseen events.

4.4 EXPLAINABILITY AND TRUST IN INFERENCING

Information systems can be called “knowledge systems” as they are designed to support the user in acquiring knowledge. They do so by providing true statements and a justification for believing these statements. In the simplest case, a knowledge system is accepted as an authority, so that the fact that it provides a certain information is enough evidence to believe it: “I know that *p* because the system said so” just like “I know that it is two o’clock because that’s what my watch says.” However, if the system is not sufficiently reliable to count as an authority, then it has to provide further evidence to support its statements, e.g.: “*p* is the case, because *q* is the case and whenever *q* then *p*.” Further information can also be required by the user to properly interpret the system’s output. For example: for turning the diagnosis of a medical information system into action – e.g., for choosing the right medication – further information on the reasons for the diagnosis might be needed. In the case of autonomous agents acting as proxies for humans in dangerous situations, explainability and trust are essential for successful collaboration.

One approach is to make an information system transparent (glass box technology) and give the user an insight into the actual generation of a statement, e.g., by presenting a complete deduction. With a complex system, however, this might not contribute to making the output clearer for a human user – it can easily be too hard to understand. Another approach is, therefore, to set up a dedicated component for generating understandable explanations. This component can aim at formally proving a statement, it can find supporting facts via abductive reasoning [68], or it can refer to causal models for determining potential causes of a stated fact (see Section 4.3), among other methods. Apart from providing evidence, an explanation component can detect knowledge gaps, that is, open questions which ought to be answered in order to evaluate a given statement: “The system said p. A good reason for p would be q. Check whether q is the case.” An explanation component can be a system on its own, which provides evidence for the output of arbitrary systems. A rule-base explanation component, e.g., could serve to justify the output of a Deep Neural Network (DNN). The justifications of the DNN output would not make the DNN transparent and it would have nothing to do with the actual generation of this output. Still, it could help the user to estimate its plausibility.

Metaphorically, information systems can also be named “knowledge systems” insofar as they “know” the information they have. To this end, they must have true information and there must be some form of quality management to assure that the information they possess is not true by accident but can be relied on. (The quality management provides the system-internal justification of its information state.) It is a challenge to elaborate how a “knowing” system in that sense can deal with uncertainty. A “knowing system” would be trustworthy by definition, because it would be an authority regarding its information.

We must accept that mistakes can occur. If we use the system wrongly then we may get the wrong answer. In order to recognise a false result we should be able to roughly estimate its plausibility and be able to recognise what cannot be true. We may also seek to repeat the use of the system to see if the results are identical. We may also seek to use a different (independent) system to confirm our result, using ensembles of various systems to compare or aggregate their outputs [69].

4.5 SUMMARY, OUTLOOK AND OPEN CHALLENGES

In the current themes of research discussed in this section, it becomes clear that to collaborate effectively with humans, autonomous systems must use KRR in order to usefully interpret incoming streams of sensory data. Much of this research overlaps with NLP research areas, as language is a human's tool for categorising and labelling the world around them into discrete and understandable concepts. As Saussure [70] put it:

> *Psychologically our thought – apart from its expression in words – is only a shapeless and indistinct mass. Philosophers and linguists have always agreed in recognizing that without the help of signs we would be unable to make a clear-cut, consistent distinction between two ideas. Without language, thought is a vague, uncharted nebula. There are no pre-existing ideas, and nothing is distinct before the appearance of language.*

Similarly, one can begin to see how the shapeless and indistinct mass of input that a computational system receives can be organised into increasingly sophisticated categories. This also points to the advantages of recent research combining neural network approaches that draw upon lower-level features (such as pixels and words) with symbolic approaches that are able to introduce a higher conceptual layer – abstracting over categories of objects and events in the way that the human mind does through experience of the world to draw conclusions, understand causation and entailments, and explain this reasoning. The ongoing research themes discussed here represent important steps towards a system that can process multi-modal streams of data and make human-like inferences from that information, contributing greatly to quickly and efficiently gaining situational awareness in a tactical environment and supporting effective decision-making.

For knowledge graphs at scale in particular, a variety of open challenges have been identified in the literature. Groth, Harmelen et al. [71] describe specific 'grand challenges' ranging from those being tackled today to broader challenges focusing on the role of knowledge graphs in future society. Amongst the 14 challenges identified are the needs for having different forms of knowledge representation that allows ambiguous, incomplete and erroneous knowledge to be captured; having methods which can be combined to make use of the best of symbolic and sub-symbolic approaches; and the need to be able to work across multiple knowledge graphs – accepting that different formats and semantics will always exist.

There is also the challenge of refining knowledge graph content once a graph has been created. Paulheim [72] summarises the approaches available to do this and concludes there remain challenges, in particular that current methods are often specific to one type of knowledge graph. A related challenge is being able to understand the completeness of knowledge graphs – a non-trivial task as it requires knowledge of a hypothetical knowledge graph that contains all the elements the knowledge graph should represent [73], [74].

Finally, the fields drawing upon KRR are varied and draw upon different types of backgrounds and expertise, from engineering to cognitive science. Thus, an additional challenge comes in the form of developing the expertise needed for implementation of KRR techniques, described in more detail in the section to follow.

4.6 REFERENCES

[1] Toutanova, K., Manning, C. and Singer, Y. (2003). Feature-Rich Part-of-Speech Tagging with a Cyclic Dependency Network. Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology 1.

[2] Tratz, S. and Hovy, E. (2011). A Fast, Accurate, Non-Projective, Semantically-Enriched Parser. Conference on Empirical Methods in Natural Language Processing. Edinburgh.

[3] Gildea, D. and Jurafsky, D. (2002). Automatic Labeling of Semantic Roles. Computational Linguistics.

[4] Palmer, M., Gildea, D. and Kingsbury, P. (2005). The Proposition Bank: An Annotated Corpus of Semantic Roles.

[5] Song, A., Bies, A., Strassel, S., Riese, T., Mott, J., Ellis, J., Wright, J., Kulick, S., Ryant, N. and Ma, X. (2015). From Light to Rich ERE: Annotation of Entities, Relations, and Events. The 3rd Workshop on EVENTS: Definition, Detection, Coreference, and Representation. Denver, Colorado, Association for Computational Linguistics.

[6] Kipper, K., Korhonen, A., Ryant, N. and Palmer, M. (2008). A Large-Scale Classification of English Verbs. Language Resources and Evaluation 42(1), pp. 21-40.

[7] Baker, C.F., Fillmore, C.J. and Lowe, J.B. (1998). The Berkeley FrameNet Project.

[8] Ding, L. and Daniel, G. (2010). Semantic Role Features for Machine Translation.

[9] Firth, J.R. (1957). Papers in Linguistics 1934 – 1951, Oxford University Press.

[10] Salton, G., Wong, A. and Yang, C.S. (1975). A Vector Space Model for Automatic Indexing.

[11] Ripley, B.D. (1996). Pattern Recognition and Neural Networks, Cambridge University Press.

[12] Mikolov, T., Chen, K., Corrado, G. and Dean, J. (2013). Efficient Estimation of Word Representations in Vector Space.

[13] Vylomova, E., Rimell, L., Cohn, T. and Baldwin, T. (2015). Take and Took, Gaggle and Goose, Book and Read: Evaluating the Utility of Vector Differences for Lexical Relation Learning.

[14] Xu, C., Bai, Y., Bian, J., Gao, B., Wang, G., Liu, X. and Liu, T.-Y. (2014). RC-NET: A General Framework for Incorporating Knowledge into Word Representations. United States, North America, ACM Press.

[15] Mikolov, T., Yih, W.-t. and Zweig, G. (2013). Linguistic Regularities in Continuous Space Word Representations. 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Atlanta, Georgia.

[16] Faruqui, M., Dodge, J., Jauhar, S.K. Dyer, C., Hovy, E. and Smith, N.A. (2014). Retrofitting Word Vectors to Semantic Lexicons.

[17] Ganitkevitch, J., Van Durme, B. and Callison-Birch, C. (2013). PPDB: The Paraphrase Database. Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies.

[18] Fellbaum, C. (ed.) (1998). WordNet: An Electronic Lexical Database, MIT Press.

[19] Levy, O. and Goldberg, Y. (2014). Dependency-Based Word Embeddings. 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). Baltimore, Maryland.

[20] Soleimani, A. and Nasrabadi, N.M. (2018). Convolutional Neural Networks for Aerial Multi-Label Pedestrian Detection. 2018 21st International Conference on Information Fusion (FUSION), pp. 1005-1010, doi: 10.23919/ICIF.2018.8455494.

[21] Karpathy, A., Joulin, A. and Fei-Fei, L. (2014). Deep Fragment Embeddings for Bidirectional Image Sentence Mapping. Advances in Neural Information Processing Systems 27, pp. 1889-1897.

[22] Socher, R., Karpathy, A., Le, Q.V., Manning, C.D. and Ng, A.Y. (2014). Grounded Compositional Semantics for Finding and Describing Images with Sentences. Transactions of the Association for Computational Linguistics.

[23] Deng, J., Dong, W., Socher, R., Li, L., Li, K., and Li, F. (2009). ImageNet: A Large-Scale Hierarchical Image Database, IEEE, pp. 248-255.

[24] Joslyn, C., Robinson, M., Smart, J., Agarwal, K., Bridgeland, D., Brown, A., Choudhury, S., Jefferson, B. Praggastis, B., Purvine, E., Smith, W.P. and Zarzhitsky, D. (2018). HyperThesis: Topological Hypothesis Management in a Hypergraph Knowledgebase. NIST Text Analytics Conference (TAC).

[25] Bangalore, S., Hakkani-Tür, D. and Tur, G. (2006). Introduction to the Special Issue on Spoken Language Understanding in Conversational Systems. Speech Communication 48(3), pp. 233-238.

[26] Issar, S. and Ward, W. (1993). CMLPs Robust Spoken Language Understanding System. Third European Conference on Speech Communication and Technology. Berlin, Germany.

[27] Chen, Y.-N., Hakkani-Tür, D. and Tur, G. and Gao, J. (2016). End-to-End Memory Networks with Knowledge Carryover for Multi-Turn Spoken Language Understanding. The 17th Annual Meeting of the International Speech Communication Association (INTERSPEECH 2016). San Francisco.

[28] Hakkani-Tür, D. and Tur, G., Asli, C. and Chen, Y.-N. (2016). Multi-Domain Joint Semantic Frame Parsing using Bi-directional RNN-LSTM. The 17th Annual Meeting of the International Speech Communication Association (INTERSPEECH 2016). San Francisco.

[29] Chai, J.Y., Fang, R., Liu, C. and She, L. (2017). Collaborative Language Grounding Toward Situated Human-Robot Dialogue.

[30] Schubert, L.K. (2015). Semantic Representation. 29th AAAI Conference (AAAI15). Austin, TX.

[31] Steedman, M. and Baldridge, J. (2009). Combinatory Categorical Grammar. Nontransformational Syntax: A Guide to Current Models. Oxford, Blackwell.

[32] Banarescu, L., Bonial, C., Cai, S., Georgescu, M., Griffitt, K., Hermjakob, U., Knight, K., Koehn, P., Palmer, M., and Schneider, N. (2013). Abstract Meaning Representation for Sembanking.

[33] Chai, J.Y., Lanbo, S., Rui, F., Ottarson, S., Littley, C., Changsong, L. and Hanson, K. (2014). Collaborative Effort Towards Common Ground in Situated Human-Robot Dialogue, ACM, pp. 33-40.

[34] Bonial, C., Donatelli, L., Abrams, M., Lukin, S.M., Tratz, S., Marge, M., Artstein, R., Traum, D. and Voss, C. (2020). Dialogue-AMR: Abstract Meaning Representation for Dialogue. 12th Language Resources and Evaluation Conference. Marseille, France.

[35] Abrams, M., Bonial, C. and Donatelli, L. (2020). Graph-to-Graph Meaning Representation Transformations for Human-Robot Dialogue. Proceedings of the Society for Computation in Linguistics 2020. New York, Association for Computational Linguistics, pp. 250-253.

[36] Austin, J.L. (1975). How to Do Things with Words, Oxford University Press.

[37] Bruce, B.C. (1978). Generation as a Social Action. Theoretical Issues in Natural Language Processing.

[38] Cohen, P.R. and Perrault, C.R. (1979). Elements of a Plan-Based Theory of Speech Acts.

[39] Allen, J.F. and C. R. Perrault (1980). Analyzing Intention in Utterances.

[40] Traum, D. (1999). Speech Acts for Dialogue Agents.

[41] Searle, J.R. (1969). Speech Acts; An Essay in the Philosophy of Language.

[42] Bunt, H., Alex, J., Choe, J.-W., Chengyu Fang, A., Hasida, K., Petukhova, V., Popescu-Belis, A. and Traum, D. (2012). ISO 24617-2: A Semantically-Based Standard for Dialogue Annotation.

[43] Marge, M., Bonial, C., Foots, A., Hayes, C., Henry, C., Pollard, K., Artstein, R., Voss, C. and Traum, D. (2017). Exploring Variation of Natural Human Commands to a Robot in a Collaborative Navigation Task. First Workshop on Language Grounding for Robotics. Vancouver, Canada, Association for Computational Linguistics.

[44] Bastianelli, E., Giuseppe, C., Danilo, C., Iocchi, L., Roberto, B. and Nardi, D. (2014). HuRIC: A Human Robot Interaction Corpus. Italy, Europe.

[45] Fillmore, C.J. (1985). Frames and the Semantics of Understanding. Quaderni di Semantica 6(2), pp. 222-254.

[46] Shen, H. (2018). Semantic Parsing in Spoken Language Understanding Using Abstract Meaning Representation. United States, North America, Brandeis University.

[47] Allwood, J., Nivre, J. and Ahlsén, E. (1993). On the Semantics and Pragmatics of Linguistic Feedback.

[48] Bunt, H. (2005). A Framework for Dialogue Act Specification. SIGSEM WG on Representation of Multimodal Semantic Information.

[49] Bunt, H. (2006). Dimensions in Dialogue Act Annotation.

[50] O'Gorman, T., Regan, M., Griffitt, K., Hermjakob, U., Knight, K. and Palmer, M. (2018). AMR Beyond the Sentence: The Multi-Sentence AMR Corpus. 27th International Conference on Computational Linguistics. Santa Fe, New Mexico, USA, pp. 3693-3702.

[51] Kim, G., Kane, B., Duong, V., Mendiratta, M., McGuire, G., Sackstein, S., Platonov, G. and Schubert, L. (2019). Generating Discourse Inferences from Unscoped Episodic Logical Formulas. First International Workshop on Designing Meaning Representations. Florence, Italy.

[52] Hajicova, E. (2019). A Plea for Information Structure as a Part of Meaning Representation. First International Workshop on Designing Meaning Representations. Florence, Italy.

[53] Lewis, D. (1974). Causation. The Journal of Philosophy 70.

[54] Hüttemann, A. (2018). Ursachen. de Gruyter.

[55] Malinas, G. and Bigleow, J. (2016). Simpson's Paradox. Stanford Encyclopedia of Philosophy.

[56] Pearl, J., Glymour, M. and Jewell, N.P. (2016). Causal Inference in Statistics: A Primer. Switzerland, Europe, John Wiley & Sons.

[57] Hitchcock, C. (2019). Causal Models. Stanford Encyclopedia of Philosophy.

[58] Pearl, J. (2019). The Seven Tools of Causal Inference, With Reflections on Machine Learning. Communications of the ACM Febuary.

[59] Pearl, J. (2009). Causality. Models, Reasoning, and Inference, Cambridge University Press.

[60] Tandon, N., de Melo, G. and Weikum, G. (2017). WebChild 2.0: Fine-Grained Commonsense Knowledge Distillation. Annual Meeting of the ACL, p. 115.

[61] Hobbs, J. (2005). Toward a Useful Concept of Causality for Lexical Semantics. Journal of Semantics 22(2), pp. 181-209.

[62] Schuler, K.K. (2005). VerbNet: A Broad-Coverage, Comprehensive Verb Lexicon.

[63] Dunietz, J., Levin, L. and Carbonell, J. (2017). The BECauSE Corpus 2.0: Annotating Causality and Overlapping Relations. 11th Linguistic Annotation Workshop. Valencia, Spain, pp. 95-104.

[64] Pustejovsky, J., Castaño, J., Ingria, R., Saurí, R., Gaizauskas, R., Setzer, A., Katz, G., and Radev, D.R. (2003). TimeML: Robust Specification of Event and Temporal Expressions in Text.

[65] Mostafazadeh, N., Grealish, A., Chambers, N., Allen, J. and Vanderwende, L. (2016). CaTeRS: Causal and Temporal Relation Scheme for Semantic Annotation of Event Structures. Fourth Workshop on Events. San Diego, California, pp. 51-61.

[66] Gordon, J. and Van Durme, B.D. (2013). Reporting Bias and Knowledge Acquisition. United States, North America, Association for Computing Machinery.

[67] Sap, M., LeBras, R., Allaway, E., Bhagavatula, C., Lourie, N., Rashkin, H., Roof, B., Smith, N.A. and Choi, Y. (2018). ATOMIC: An Atlas of Machine Commonsense for If-Then Reasoning.

[68] Brachman, R. and Levesque, H. (2004). Knowledge Representation and Reasoning. Morgan Kaufmann.

[69] Zhou, Z. (2012). Ensemble Methods: Foundations and Algorithms, Chapman & Hall/CRC.

[70] Saussure, F.M. (2011). Course in General Linguistics, Columbia University Press.

[71] Groth, P., van Harmelen, F., Ngomo, A.-C.N., Presutti, V., Sequeda, J. and Dumontier, M. (2018). Grand Challenges. Knowledge Graphs: New Directions for Knowledge.

[72] Paulheim, H. (2016). Knowledge Graph Refinement: A Survey of Approaches and Evaluation Methods. Semantic Web.

[73] Darari, F., Nutt, W., Pirrò, G. and Razniewski, S. (2018). Completeness Management for RDF Data Sources. ACM Transactions on the Web 12(3).

[74] Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., Melo, G.D. Gutierrez, C., Gayo, J.E.L., Kirrane, S. Neumaier, S., Polleres, A., Navigli, R., Ngomo, A.-C.N., Rashid, S.M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S. and Zimmermann, A. (2020). Knowledge Graphs. arXiv.

Chapter 5 – CONCLUSIONS AND RECOMMENDATIONS FOR FUTURE TECHNICAL PROPOSALS

As we make substantial progress at sharing and integrating data across the alliance, we are being confronted with the very real risk of data paralysis. Whether we have truly reached the point of drowning in our own data, we must be prepared for that eventuality. As stated earlier, the attention of human analysts is still likely to be overstretched by the need to connect subtle but significant observations across multiple domains and KRR is an important part of the solution to address this issue in a timely manner.

Our introduction to this report focused on the increasing challenge, for NATO capabilities and national forces, of making sense of an increasingly complex range of data sources, alongside the increasing complexity of the operational challenge. In order to achieve a state of 'Information Advantage', in particular in achieving high level fusion in support of military decision making we observe that KRR methods could play a central role – the foundation assumption for establishing IST-ET-111.

In the course of IST-ET-111 we have reviewed the current state of the art in KRR and sought to understand the opportunities and challenges of KRR. We have reviewed what KRR is, considered the implementation of KRR and reviewed some of key research themes. In this section we draw this together with some conclusions and offer recommendations for future Technical Activities under the umbrella of the NATO STO.

5.1 CONCLUSIONS – THE OPPORTUNITIES OF KRR METHODS

So, can KRR help us addresses the challenges of attention in an information rich world and deal with the complexity posed by the current Defence challenge?

We propose that KRR offers working solutions to some of today's challenges, underpinned by more than 40 years of research. However, we note that this research has not been fully exploited in military capabilities, with specific challenges in the process of the 'knowledge engineering' required to capture and record domain knowledge in a suitable form.

Nevertheless, in recent years there are examples of mature systems for knowledge representation (such as the MIP) and for new initiatives (for example DICO) building on the standards which underpin the aspirations for the semantic web, and on fundamental frameworks such as BFO. **As such initiatives gain traction there will be an opportunity to make increasing use of sophisticated information and knowledge techniques, including the inferencing of new knowledge to support operational decision making.** It is also likely that the increasing interest and take-up of knowledge graphs by key technology players (Google, Facebook, eBay, Microsoft etc) will provide future opportunities and motivate further interest in such approaches.

A key opportunity, identified by IST-ET-111, is the ability for the current wave of ML systems to be complemented by KRR methods in order to improve the explainability of results and, therefore, developing trust in future AI systems, while also allowing AI systems to cope with situations with minimal training data. However, to date, we recognise work to combine the symbolic (knowledge-based) methods and sub-symbolic (computational approaches based on training data) is in an early stage. And making the case for such future development has to be done against defence's enthusiasm for the 'low hanging fruit' currently offered by sub-symbolic methods, driven by data availability, computational power, access to tooling/computational frameworks and success in other domains. **We believe, however, there is an opportunity to drive early work to develop the tools, skills and relationships required for NATO nations to exploit knowledge-based (symbolic) systems alongside the computational capabilities offered by the current crop of ML (sub-symbolic) systems.**

As an underpinning theme KRR has a place in almost any technical activity related to Data, AI or Autonomy. In considering the way ahead, we note there are many related activities under the umbrella of the NATO STO. This is a complex picture. In particular we highlight alignment with the following groups, while offering a longer list of related activates in Annex A:

- IST-165 – High Level Fusion of Hard and Soft Information for Intelligence.
- IST-ET-112 – ML Ecosystem.
- IST-157 – Human in the Loop Considerations for AI.
- IST-173 – Mission-Oriented Research for AI and Big Data for Military Decision Making.

To avoid duplication of effort we propose any future activity be closely coordinated with appropriate technical activities.

5.2 CONCLUSIONS – THE NEED FOR UNDERPINNING SKILLS AND EXPERTISE

We have recognised that implementing knowledge representation methods can be complex, and that to mitigate the temptation to “model the world” any knowledge engineering activity must retain a clear sense of purpose and adopt modular approaches, for example by defining higher level ontologies to model general concepts, while lower level ontologies are used to model domain specific concepts. We have described how the W3C Semantic Web stack offers an accessible approach that allows knowledge representation activities to build on common standards, supporting such modular approaches. **These standards, and the expertise in specifying systems against them and developing capabilities with them, will become important skillsets for the future NATO workforce if we are to achieve goals for high level fusion.**

Indeed, there are wider needs around skills and expertise if NATO is to exploit effective KRR. **The ’knowledge engineering’ skill set is critical and may now support the need to recognise formally a ‘knowledge engineer’ role.** Knowledge engineers require knowledge and experience of KRR methods and the other technologies highlighted in this report, but they also require the skills to interact with subject matter experts with unstructured and structured techniques to allow them to dissect and represent a particular domain in the most appropriate manner. **Such skills do not ‘grow on trees’ and might be best developed collaboratively across any nascent capability in nations, a good first step being to widen awareness and undertake a stocktake of the technical capability already available.**

We have also described how new conceptual approaches to KRR might offer opportunities to develop different analytical approaches, for example adopting techniques such as 4D ontologies that focus on the states of entities and support more expressive description and analysis.

5.3 CONCLUSIONS – CURRENT RESEARCH THEMES

The work of IST-ET-111 has not focused on any one domain, although many of the examples offered are focused on command and control and intelligence analysis use cases. However, we have also identified opportunities to use KRR techniques to support human-machine interaction techniques which will be increasingly important as Defence seeks to implement human-machine teams and sophisticated autonomous systems. **Indeed, while KRR methods in support of analysis and high level fusion goals can undoubtedly offer future capability improvements, it is to the future realisation of effective human-machine teams that KRR may be able to offer significant opportunities.** This also includes how improved explainability could improve trust in machine analysis and future autonomous decision making.

There are also a number of areas where ongoing research could develop further that which can already be achieved. While natural language processing has made significant steps in recent years **there are opportunities for defence capabilities to use this work to exploit large quantities of free text material in knowledge bases/graphs, and for the knowledge base/graph itself to be used to improve that analysis.**

Other research is focusing on understanding how true causal relationships can be established between elements of a knowledge base/graph. Recognising that the current approach in sub-symbolic methods is focused on identifying correlation in data **exploiting developing thinking in causal modelling could improve robust decision support and further draw together links between symbolic and sub-symbolic methods.**

5.4 RECOMMENDATIONS

Against these conclusions we offer the following outline Technical Activity Proposals for future work under the NATO STO.

We concluded that KRR (symbolic) methods are mature, but have previously been limited by the needs of knowledge engineering for a particular domain. However, as the current rise of ML (sub-symbolic) methods illustrate their high demand for data, difficulty in identifying rare events and the user's desire for increased trust in their results we believe there is a significant opportunity to explore the complementary aspects of the two approaches.

Recommendation 1: The NATO STO sponsors a technical activity to demonstrate the complementary use of symbolic and sub-symbolic methods and their benefit to improved decision making.

Such an activity could help to focus expertise and capability across the alliance. The work should consider the design of hybrid approaches to sense-making, the exploitation of large volume of free text (as knowledge stores), how the completeness of knowledge graphs can be estimated, and how such hybrid systems support improved explainability.

We note that any activity relating to Recommendation 1 should be aligned or joined with suitable existing activities to avoid stovepipes and to ensure the optimum approach to any burden sharing. For example, in the course of IST-ET-111 we have not considered the availability of datasets to drive experimentation, development and benchmark commercial capabilities, but this requires further consideration if we wish to drive future work to explore the complementary nature of symbolic and sub-symbolic methods for defence applications. Such dataset generation is non-trivial and should be considered in any work arising from Recommendation 1.

Furthermore, we noted that wider adoption of such hybrid systems will require the development of new skills and expertise in knowledge engineering, and the adoption of key standards (which are largely already in existence). IST-ET-111 represents a nascent group of expertise from which further capability development could be achieved.

Recommendation 2: The NATO STO sponsors a suitable technical activity to support a virtual lecture series/workshop to increase the awareness of KRR technologies in the science and operational sectors of NATO, in order to provide a catalysis for further skills development in this area.

Finally, a significant opportunity exists for the use of KRR methods to support future human-machine teams, and key research themes, such as the use of knowledge to constraint text analytics and the definition of causal models will develop this further. Recognising the relative maturity of text analytics (and that it represents a tangible use case for work under Recommendation 1) IST-ET-111 recommends on-going exploratory activity.

Recommendation 3: The NATO STO sponsors a dedicated Exploratory Team to consider specific interests in causal modelling and its application to knowledge-based systems, as a possible precursor to future practical demonstrations under activities such as that against Recommendation 1.

As an area of development, the team consider causality and causal modelling represents a significant opportunity to improve decision making, although we recognise that further developments are required. Further work exploring this area, including developing relationships with the academic sector, would allow future work to rapidly exploit opportunities as they become available.

Annex A – RELATED NATO STO ACTIVITIES

The following activities have links to the themes of IST-ET-111:

- AVT-ET-204 Data Fusion and Assimilation for Scientific Sensing and Computing.
- HFM-178 Meaningful Human Control (MHC) Over AI-Based Systems.
- HFM-322 Meaningful Human Control of AI-based Systems: Key Characteristics, Influencing Factors and Design Considerations.
- IST-157 Human in the Loop Considerations for Artificial Intelligence.
- IST-165 (AI2S) High-Level Fusion of Hard and Soft Information for Intelligence.
- IST-169 Robustness and Accountability in Machine Learning Systems.
- IST-173 (AI2S) Mission-Oriented Research for AI and Big Data for Military Decision Making.
- IST-177 Social Media Exploitation for Operations in the Information Environment.
- IST-178 Big Data Challenges: Situation Awareness and Decision Support.
- IST-ET-112 ML Ecosystem.
- SAS-157 Automation in the Intelligence Cycle.
- SCI-331 Fostering and Managing the STO Autonomy Portfolio.
- SET-263 Swarms Systems for Intelligence Surveillance and Reconnaissance.
- SET-278 Machine Learning for Wide Area Surveillance.
- SET-279 Space-Based Radar Systems, Big Data and Artificial Intelligence.
- SET-283 Advanced Machine Learning ATR Using SAR/ISAR Data.

Annex B – MIP INFORMATION MODEL AND RICH EVENT ONTOLOGY

B.1.1 MIP INFORMATION MODEL (MIM)

The MIP Information Model (MIM) defines common semantics for the Command & Control (C2) domain. It adopts and consolidates concepts from various authoritative sources, mostly NATO standards. The main objective of the MIM is to support information exchange in joint and combined operations. Custodian of the MIM is the Multilateral Interoperability Programme (MIP), a military standardisation body comprising 24 member nations, NATO, and EDA (European Defence Agency).

The core of the MIM is a taxonomy with thousands of militarily relevant concepts. This includes basic battlespace concepts – objects and actions – as well as staff concepts such as plans, overlays, and organisation structures. The set of supported objects includes persons and organisations, equipment and consumables, facilities, as well a control features (e.g., organisation boundaries) and meteorological/geographical conditions. The concepts in the MIM have a rich set of properties to describe the inherent characteristics, the status and the capabilities of their instances as well as their relationships with each other. In addition, metadata can specify, e.g., the originator, the confidentiality, the validity period, and the appraisal of exchanged information. All elements are provided with an operational definition to support a common understanding. In total, the MIM, version 5.1 comprises >900 classes/data types, >1200 attributes, >170 associations, <500 enumerations and <6900 literals. The number of concepts (classes + MIM-specific category codes) is <3200.

The MIM is specified in the Unified Modelling Language (UML). It is platform-independent, i.e., it is not tied to a specific exchange technology. By means of UML stereotypes, the elements of the MIM are semantically "annotated". For instance, the meaning and intended use of attributes is further specified by stereotypes such as "name" and "measure", which introduce additional information, e.g., the unit of measure.

Communities of Interest (COIs) can adopt the MIM for developing interoperability specifications in support of their specific processes. The MIM comes with a tool suite that allows to subset/extend the model and to define structured messages, which are structurally compliant with the MIM. Model transformations allow the generation of representations in XML Schema and OWL.

All MIM-related products can be found at https://www.mimworld.org. The website hosts introductory documentation, the information model in Sparx Enterprise Architect format, services to browse the model online, and the tool suite. For general information on the Multilateral Interoperability Programme (MIP), please see https://www.mip-interop.org.

B.2.1 RICH EVENT ONTOLOGY (REO) – ONTOLOGICAL HUB FOR EVENT REPRESENTATIONS

There are a variety of valuable ontological resources that focus on representing world knowledge, in the philosophical tradition of ontology, and may or may not include lexical resources mapping to those concepts. For example, both the Basic Formal Ontology (BFO) described by Smith et al. [1] and the Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) described by Masolo et al. [2] and Aldo et al. [3] focusses on modelling the concepts underlying natural language and common sense, as opposed to the language itself. In contrast, WordNet [4], UBY [5] and the Ontologies for Linguistic Annotation (OLiA; [6]) focus on representing lexical and morpho-syntactic information in an ontology.

The Rich Event Ontology (REO) aims to represent lexical event semantic information; thus, it hopes to capture both common-sense world knowledge about events and their participants, as well as lexical information on how these concepts are realised and tagged in various English annotation schemas [7], [8].

REO unifies existing Semantic Role Labelling (SRL) schemas used in Natural Language Processing (NLP) by providing an independent conceptual backbone through which they can be associated, and it augments the schemas with event-to-event causal and temporal relations. Specifically, REO brings together the SRL labelling schemas of FrameNet [9], the Rich Entities Relations and Events (ERE) project [10] (originally based on Automatic Content Extraction (ACE) [11], and VerbNet [12]. FrameNet, ERE, and VerbNet have wide-coverage lexicons of events, and they contribute annotated corpora and additional semantic and syntactic information that can be crucial to identifying events and their participants. REO serves as a shared hub for the disparate annotation schemas and therefore enables the combination of SRL training data into a larger, more diverse corpus, as well as expanding the set of lexical items associated with each event type. By adding temporal and causal relational information, the ontology also facilitates reasoning on and across documents, revealing relationships between events that come together in temporal and causal chains to build more complex scenarios. Most recently, REO has been enriched with Generative Lexion "qualia relations" [13] which specify how an entity is used, how it was created, its component parts, and its formal type [14]. REO remains under development and is being leveraged in experimentation as part of the world model for autonomous agents in search and navigation tasks [15].

B.3.1 REFERENCES

[1] Smith, B., Almeida, M., Bona, J., Brochhausen, M., Ceusters, W., Courtot, M., Dipert, R., Goldfain, A., Grenon, P. and Hastings, J. (2014). Basic Formal Ontology 2.0 Draft Specification and User's Guide.

[2] Masolo, C., Borgo, S., Gangemi, A., Guarino, N. and Oltramari. A. (2003). WonderWeb Deliverable 18:Ontology Library (Final). http://wonderweb.man.ac.uk/deliverables/D18.shtml.

[3] Aldo, G., Nicola, G., Claudio, M., Alessandro, O. and Luc, S. (2002). Sweetening Ontologies with DOLCE.

[4] Fellbaum, C. (ed.) (1998). WordNet: An Electronic Lexical Database, MIT Press.

[5] Gurevych, I., Eckle-Kohler, J., Hartmann, S., Matuschek, M., Meyer, C.M. and Wirth, C. (2012). UBY – A Large-Scale Unified Lexical-Semantic Resource Based on LMF. Germany, Europe.

[6] Christian, C. and Sukhareva, M. (2015). OLiA – Ontologies of Linguistic Annotation. Semantic Web 6(4).

[7] Bonial, C., Tahmoush, D., Brown, S.W. and Palmer, M. (2016). Multimodal Use of an Upper-Level Event Ontology. Fourth Workshop on Events. San Diego, California, Association for Computational Linguistics, pp. 18-26.

[8] Brown, S., Bonial, C., Obrst, L. and Palmer, M. (2017). The Rich Event Ontology. NEWS@ACL 2017, pp. 87-97.

[9] Fillmore, C.J., Baker, C.F., Sato, H. (2002). The FrameNet Database and Software Tools. Proceedings of the Third International Conference on Language Resources and Evaluation (LREC'02), Las Palmas, Canary Islands, Spain, May 2002.

[10] Song, A., Bies, A., Strassel, S., Riese, T., Mott, J., Ellis, J., Wright, J., Kulick, S., Ryant, N. and Ma, X. (2015). From Light to Rich ERE: Annotation of Entities, Relations, and Events. The 3rd Workshop on EVENTS: Definition, Detection, Coreference, and Representation. Denver, Colorado, Association for Computational Linguistics.

[11] Doddington, G., Mitchell, A., Przybocki, M., Ramshaw, L., Strassel, S. and Weischedel, R. (2004). The Automatic Content Extraction (ACE) Program Tasks, Data, and Evaluation. Fourth International Conference on Language Resources and Evaluation (LREC'04). Lisbon, Portugal, European Language Resources Association (ELRA).

[12] Kipper, K., Korhonen, A., Ryant, N. and Palmer, M. (2008). A Large-Scale Classification of English Verbs. Language Resources and Evaluation 42(1), pp. 21-40.

[13] Pustejovsky, J., Havasi, C., Littman, J., Rumshisky, A. and Verhagen, M. (2006). Towards a Generative Lexical Resource: The Brandeis Semantic Ontology.

[14] Kazeminejad, G., Bonial, C., Brown, S.W. and Palmer, M. (2018). Automatically Extracting Qualia Relations for the Rich Event Ontology. COLING: International Conference on Computational Linguistics, p. 2644.

[15] Lukin, S.M., Bonial, C. and Voss, C.R. (2019). Visual Understanding and Narration: A Deeper Understanding and Explanation of Visual Scenes.

Annex C – DEFENSE INTELLIGENCE CORE ONTOLOGY (DICO)

The DICO is the Defense All-source Analytic Enterprise (DIAAE) knowledge model for Object-Based Production (OBP). It provides the semantic framework to access and organise defence intelligence data in a way that is intuitive and mission focused for the DIAAE analysts and collectors preparing for, and participating in, dynamic conflicts. DICO is a mid-level ontology, designed and built according to Basic Formal Ontology (BFO) standards. Within the DICO, concepts are modelled using real world relationships that are structured in a way that is meaningful to both computers and humans.

This approach enables standards-based information exchange and interoperability between integrated applications and services. The expressivity and flexibility of the DICO knowledge model allows for future development, information sharing, and analytics. The DICO will facilitate the:

- Consistent development of classes and relationships that reflect the content found in authoritative Defense Intelligence Analytic Program (DIAP) sources such as the Modernized Integrated Database (MIDB);
- Ability to better incorporate spatio-temporal entities (e.g., the movement of mobile missiles out of garrison) with current and future analytic tool suites and data bases focused on fixed entities such as facilities;
- Enhanced (i.e., computer-assisted tools such as ML) reasoning that supports intelligence analysis methods instead of data dictating analysis;
- Integration of relevant data from disparate intelligence sources and publicly available sources into a common object management service;
- Logical and consistent expansion of reasoning support into any domain at any level of granularity (i.e., from large aggregate objects down to elemental parts of objects); and the
- Improved usage of Intelligence Functional Codes (IFCs) and other Intelligence Community coding systems to reason with intelligence and analyse production.

The primary form of implementation of the DICO will be integration with the Object Management Service (OMS) of the Defense Intelligence Agency (DIA). Creation and implementation of DICO compliant ontologies at the application level will strengthen semantic integration across the Defense Intelligence Enterprise and add to an evolving Enterprise Knowledge Graph. Note that the following discussion of the DICO makes use of material from several previously-authored sources [1], [2], [3], [4].

As discovered by Rohr and Miller in their report from 2016, too often, well-intended attempts are made at addressing data interoperability by developing standards for local to cross-organisational use, but the results have fallen short of meeting DIAAE-wide data interoperability and semantic understanding. Even when data standards are developed, they are often not used or are overcome by changes in technology or missions. Reasons for this lack of adoption include a lack of full awareness of data standards, data standards that are not enforced, data standards that are developed with insufficient analytic unit engagement, and an insufficient understanding of the value that data standards provide. Importantly, when the data standards are not developed with internally coherent logic, coupled with user engagement, the result will be a standard that is simply unusable. The structure of the DICO, however, will improve interoperability and semantic understanding of intelligence data used by the DIAAE. This logical, user semantics-based structure, coupled with data standards aligned to the DICO, will greatly improve potential for adoption of both the DICO and related data standards.

The DICO is modelled using Resource Description Framework (RDF) and Web Ontology Language (OWL) [5] to provide the knowledge representation language and formalism necessary to represent consistent definitions of and relations among concepts. These languages are W3C standards and adherence to these standards improves interoperability.

As mentioned above, the DICO has been developed according to BFO standards (specifically, ISO/IEC 21838-2). BFO is a small, domain neutral, upper-level ontology that will provide a common semantic framework able to integrate disparate concepts. DICO also integrates with and extends from Common Core Ontologies (CCO) where possible. CCO, constructed primarily by CUBRC for several Departments of Defence (DoD) research organisations, comprise eleven mid-level ontologies that extend directly from BFO. CCO's purpose is to represent and integrate taxonomies of generic classes and relations across several domains of interest such as material objects, time, measurements, and space.

The DICO is developed with the primary purpose of improving our understanding of and analysis within defence intelligence domains such as infrastructure, order of battle, and targeting. Therefore, the DICO extends from CCO terms where practical and adds additional terms where necessary to better model these particular domains and their interdependencies as depicted in Figure C-1.

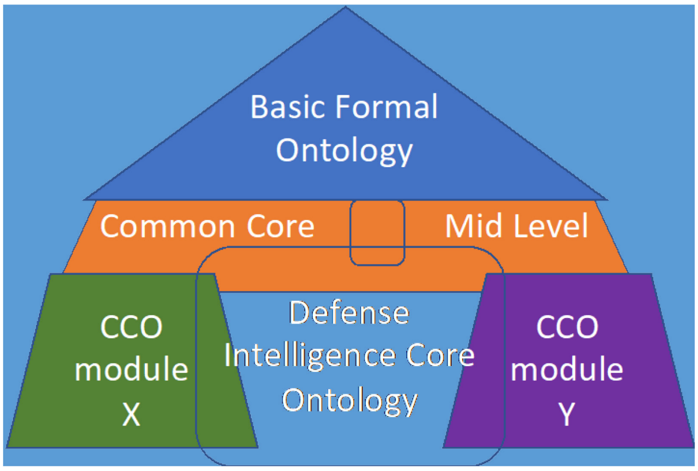

Figure C-1: DICO, BFO, and CCO.

The first versions of DICO will be necessarily less expressive than is possible following BFO and CCO modelling practices. DICO is governed by the taxonomical structure and logical reasoning inherited from BFO-CCO. However, this will not negatively constrain innovation and creativity at the application level of ontology development. Newly added content that is compliant with the BFO-CCO-DICO integrating framework will add to the overall Defense Intelligence Enterprise Knowledge Graph – an evolving, logic governed, computable graph.

C.1.1 DICO DEVELOPMENT PROCESS, DESIGN PRINCIPLES, AND BEST PRACTICES

Since domain ontologies are designed primarily to support the operators' – in this case intelligence analysts' – ability to reason with their data, it is best to start with sample competency questions that would help illustrate the types of knowledge one would want to extract from information via the ontology. In the field of defence intelligence, collection and analysis is often driven by Essential Elements of Information and Priority Intelligence Requirements (PIR). An example PIR could be: "What is Country X's military capability to conduct asymmetric operations against Country Y?"

While it may be possible to translate this question (and others) directly into queries against a database, they most often require decomposition into more specific questions. For example: "Show me all activity within a given area for the last 38 months".

Questions can then be transformed further into the form of "pseudo-code" to demonstrate the short bridge between a plain language question and a statement that is written in a logical format that can be understood by the query system, e.g.: "Which AREA OF OPERATIONS 'is site of' some MILITARY EXERCISE?" – an activity-based intelligence question. In fact, with a bit of Natural Language Processing (NLP), a plain language question can be fairly easily transformed into a machine-readable query statement. The resultant query can be placed directly to an analytic engine behind the user interface of the analyst's intelligence application.

Most importantly, the sample questions based in First Order Logic (things that are happening in the world as opposed to descriptions of those things) can help ontology development in two ways. First, they help us bound the domain of interest. In other words, if the concept is not somehow related to the competency questions, it reduces the need to cover the concept with entities in the ontology of interest. Second, if the current ontology does not identify the entities of interest in the competency questions, they provide a guide for how to include them. In short, the competency questions help to establish what's "in and what's out" of the ontology.

In the case of the DICO, for example, the authors gathered competency questions from across the defence intelligence all-source analytic enterprise to ensure all concepts of common concern to defence all-source intelligence analysts are captured in the DICO, or at a minimum, there is a "hook" provided for development of a more detailed module to address the analyst's needs.[1]

C.1.1.1 Uniquely Identifying Entities

Each entity and property in the DICO, and any modules created according to the DICO, has a unique Uniform Resource Identifier (URI) for unambiguous identification.[2] The URI is constructed from a base URI – a namespace that is unique to each ontology – and a local identifier. For example, the base URI of a tank within the DICO may be www.dia.mil/dico/tank/. The local identifier of some specific tank on the battlefield (i.e., an instance of a tank) may be: www.dia.mil/dico/tank/USCENTCOM_TNK_000001.

C.1.1.2 Ontology Entities and DICO Entity Categories

The DICO (and indeed all fully constructed ontologies using OWL) is comprised of six forms of entities; class terms, object properties, data properties, data types, annotations, and individuals (also called instances). Classes (also referred to as "Types" or "Universals") serve as the placeholder term for the collection of individual instances in the world. For example, the class 'Military Organization' ---has individual--- > '1st Tank Brigade'. At its highest level, BFO divides the class terms into continuants and occurrents. To facilitate understanding of the DICO we group the class (or objects) terms into the following five primary categories:[3]

- **Material Entities**. The things that exist in the world, e.g., FACILITY, VEHICLE, PERSON.
- **Dependent Entities**. Formalisation of the many different ways we describe the things around us and that we do for example, the speed and function of a tank.[4]

1 As of June 2020, the socialising with domain SMEs was on-going. Version 2 was published with draft terms to facilitate forward progress while meetings with SMEs continued.

2 http://www.w3.org/Addressing/URL/uri-spec.html.

3 Class terms are the primary elements that make up a well-formed taxonomical hierarchy consisting of parent and child classes – e.g., 'Weapon' is the parent class of 'Rail Gun'. When created using a common upper-level framework, taxonomies can be logically integrated and processed with computer algorithms. This insight is the lynchpin for semantic integration and governance.

4 According to BFO, the functions, roles, and dispositions of material objects are called 'realizable entities." Material Objects are the bearers of these realisable entities, which are always realised by some 'PROCESS' – e.g., 'FIRES FUNCTION is realised only by some 'FIRING PROCESS'.

- **Temporal Entities.** The things that occur. These are processes and temporal intervals (i.e., spans of time). All processes require some agent to conduct them in order to be relevant to analysis. However, the act itself is an important entity to codify independent of the agent conducting it. Examples include: TIME INTERVAL, VEHICLE MOVEMENT, and MILITARY OPERATION.
- **Spatial Entities.** This category includes both absolute locations, such as a geographic coordinate, and relative locations that are codified in BFO as SITES. SITES are one of the most abstract BFO entities but also very important to accurately describe a military location that is not a fixed facility. In a general sense, SITES are defined in relation to the boundaries that make them significant. A good example is a room. A room does not exist independent of its surrounding. It is an empty space that is only relevant because of the walls that enclose it. Additional spatial entities are SPATIAL REGIONS. These are absolute locations defined by a reference system with a common datum to all points in the system. Examples of spatial entities include: SURFACE TO AIR MISSILE SITE, AREA OF OPERATION, and VILLAGE.
- **Complex Entities.**[5] These are any combination of the above four categories of entities that are grouped together by virtue of their exhibiting some combination of characteristics that does not correspond to any universal (class). A good example is an EVENT which is a combination of some AGENT/s conducting some ACT at some LOCATION during some TIME INTERVAL.

Additionally, it is helpful to describe the relationship forms used to formally relate Material Entities, Dependent Entities, and others. The DICO also formally specifies many relationships between entities, and data elements about them. This formalisation helps both the analyst and computer to reason with the data much more than can be accomplished using legacy relational databases. Relationships between entities (objects) are called Object Properties. For example:

- PERSON (subject):
  - *bearerOf* (object property) DRIVER ROLE (object).
  - *participatesIn* (object property) VEHICLE MOVEMENT PROCESS (object).

Lastly, the DICO formally specifies the relationships between entities and their attributes that are not codified as entities themselves. For example:

- PERSON (subject):
  - *hasGender* (data property) **male** (literal value).

These types of attributes of an object are often conflated with other entities so that the concepts such as "country of ownership," "function of facility", and "object dimension" are all listed in a relational database as attributes of the physical object. This logical imprecision leads to confusion and the inability to leverage many computational resources effectively to solve analytic problems[6].

Data values, the actual data related to an instance by a data property (e.g., the specific range of an aircraft), are contained in knowledge graphs or relational databases linked to the DICO. They are not contained as part of the DICO itself. The graph database that contains defence intelligence information of use by the DIAAE is called an object management service. Figure C-2 shows how the DICO and an OMS together provide a cohesive, standardised web of data connecting various external sources, annotated data, and the knowledge model.

---

[5] In ontology terms, these are also called "defined classes."

[6] A good example of this issue are dates and geocoordinates. A tank does not "possess" a date or a geocoordinate so dates and coordinates should not be considered attributes of the tank. A person can make a leap in logic to understand the geo-coordinate is an attribute of a location at which the tank is stationed, but a computer cannot do that. However, a tank "does something" on a date and "is located" at a spot that has the geo-coordinate as an attribute. Once this is modelled correctly, the computer can understand the relationship between the tank and the geo-coordinate.

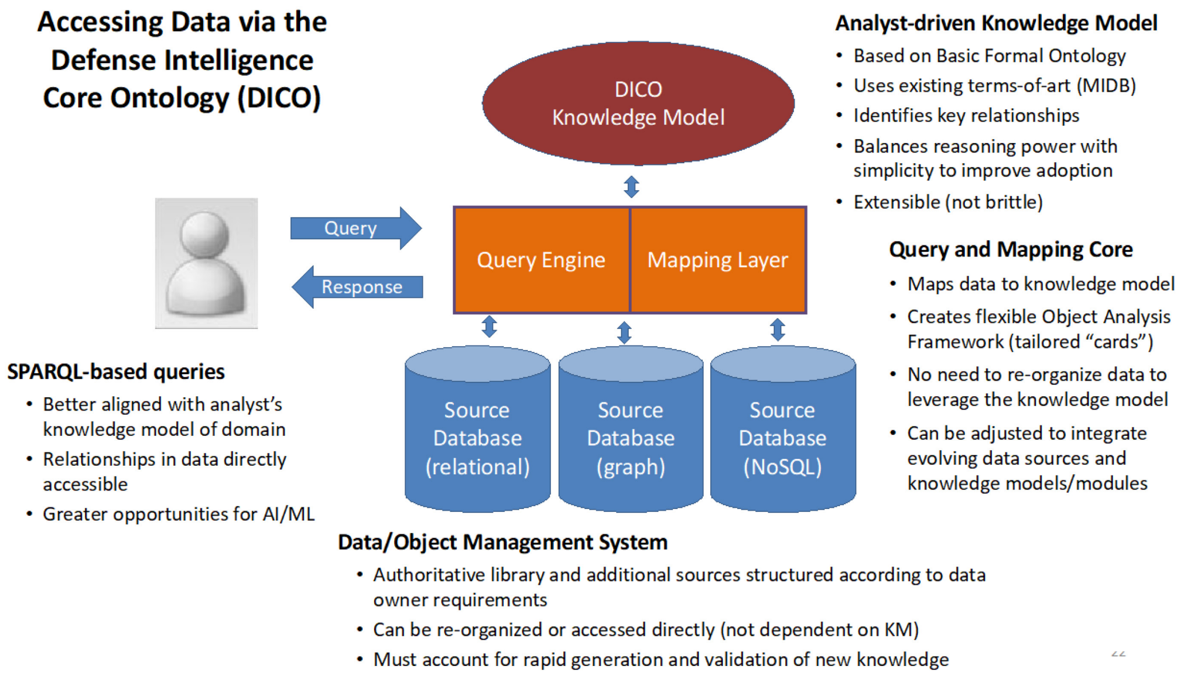

Figure C-2: Interacting with the DICO and OMS.

Importantly, an OMS graph database does not need to contain all the spatio-temporal data that is collected on each object. To do so would require a complex graph database that must be continuously updated as mobile objects move around the battlespace. Instead, many such attributes of objects in an OMS can be linked to databases purposely built for such tasks.

Lastly, successful implementation of the DICO for KRR depends on a combination of top-down and bottom-up modelling approaches to ensure alignment across all the application ontologies created according to the DICO. At DIA, the top-down piece is building the DICO from BFO and extending the model horizontally across all the domains of interest (e.g., analysis, Adversary COA development, infrastructure, order of battle, cyber operations, etc.). This ensures maximum interoperability of the models. More importantly, it ensures the formal representation of the models is also accurate and logical. After that core is created, one then proceeds bottom-up to ensure the ability to link the available data in the various domains to the knowledge model that has been carefully crafted to ensure the formal representation of the domains is accurate, logical, and fully inter-operable. Some of the process can be automated with tools such as NLP, but the initial linkage should be reviewed manually and a series of queries should be done to ensure the data is linked effectively. This top-down, bottom-up, then iterate is key to ensure effective deployment in the operational environment.

C.2.1 REFERENCES

[1] Arp, R., Smith, B., and Spear, A.D. (2015). Building Ontologies with Basic Formal Ontology, The MIT Press.

[2] Rohr, V. and Miller, S. (2016). Data Sharing and Semantic Understanding Across the Intelligence Community. El Segundo, California, Aerospace Corporation.

[3] Myers, M., Behling, J., Otte, N. and Lebo, T. (2019). Joint Intelligence Knowledge Graph Developer Guide. Johns Hopkins Applied Physics Laboratory.

[4] Rudnicki, R., Smith, B., Malyuta, T. and Mandric, W. (2016). Best Practices of Ontology Development, Cubric.

[5] W3C. (2020). Semantic Web Standards. Retrieved 17/11/2020, 2020, from https://www.w3.org/2001/sw/wiki/Main_Page.

Annex D – KNOWLEDGE REPRESENTATION AND REASONING IN PRACTICE – THE WISDOM R&D PLATFORM

D.1.1 THE WISDOM R&D PLATFORM

WISDOM is a R&D software platform [1]. It has been developed at Defence Research and Development Canada (DRDC), mainly under Project 05da: Joint Intelligence Collection and Analysis Capability (JICAC), and is meant to be a proof-of-concept prototype of an intelligence production support system. It is geared towards research in data/information/knowledge integration, fusion, analytics, management and exploitation, aiming at providing a capability to support the analysts and decision makers in developing their belief, opinion, judgment, or prediction about situations while these individuals are involved in situation analysis and decision-making activities.

WISDOM is a federation of innovative, computer-based, composable and interoperable capability units provided in the form of an inventory of web services on a service-oriented architecture, thereby facilitating system integration and interoperability. These capability units can be integrated, interleaved and composed into overall, continuous process flows for the management and exploitation of data, information and knowledge, and supporting the individuals concerned with intelligence production, enabling and facilitating the creation and maintenance of enhanced situation awareness.

WISDOM is a generic, domain agnostic platform reusable in different contexts and settings. It is composed of three main components:

- The Sense making Support System (S^3);
- The Knowledge Engineering Support System (KESS); and
- The Unified Data Space (UDaS).

The S^3 is tailored to sense making, i.e., to the process of creating situation awareness in situations of uncertainty. It provides the capability required to exploit data, information and knowledge in a way that enables and facilitates the creation and maintenance of enhanced situation awareness for the end-user. It provides a suite of integrated components to support various analytical processes such as automated reasoning/inferencing, data correlation, temporal alignment, hypothesis management, list-based processing, text analytics, and visualisation on a map, which all play a key role in the examination of a situation, its elements, and their relations, to provide and maintain a product, i.e., a state of situation awareness, for the analysts and/or decision makers.

A large portion of the current version of WISDOM makes use of knowledge-based systems technologies. Aligned with this, the KESS component of WISDOM provides a user-friendly interface that makes the definition, specification and exploitation of knowledge representation artifacts easy for the knowledge engineers and/or the end-users of WISDOM. It enables the formal encoding of domain and expert knowledge in terms of ontologies, proposition templates, propositions, graphs, situation models, spatial features, case templates, cases, inference rules, text-based templates, analysis configurations, etc., and sets of such items.

The UDS component of WISDOM is the least mature element. It is meant to be a central all-source, multi-int repository for all intelligence data, information and knowledge (sensed, observed, derived, etc.) currently available and related to the past, current and future situations of intelligence interest. Documents of all kinds can be ingested into the UDS, and documents can be retrieved from it. Among other things, the UDS must:

- Provide an all-source, multi-int data integration framework;
- Provide a unifying data framework;

- Provide a single interface to all data;
- Provide a single access point to data, information and knowledge from heterogeneous sources;
- Provide efficient search from a single point of access;
- Be able to semantically align heterogeneous sources;
- Be fully compatible with W3C Semantic Technologies;
- Preserve the traceability of information;
- Integrate and handle structured and unstructured contents;
- Provide a scalable environment to deal with huge volumes of data (once the storage and processing capabilities has been aligned on the Big Data paradigm);
- Provide a platform for data analytics;
- Provide a solid foundation for the sense making support components; and
- Enable both data pre-processing and data post-processing to support future requirements.

D.2.1 WISDOM DATA STRATEGY

A crucial aspect of the WISDOM R&D platform is the data strategy that has been developed to meet the challenging requirements for knowledge representation, as well as the demanding requirements for the persistence and exchange of data, information and knowledge. The former requirements are related to the exploitation of AI approaches and techniques (e.g., automated reasoning), while the latter are linked to the needs for interoperability between information processing services, agents and/or systems (i.e., machine-to-machine interoperability in distributed systems). DRDC has devoted significant effort to this aspect over the span of the JICAC project.

A key design decision has been made very early for WISDOM to exploit multiple representation paradigms instead of adopting a single approach that would have to fit all contexts and settings. One objective was to exploit the collective, complementary strengths of many paradigms, while avoiding their individual pitfalls and weaknesses. Another objective was to fit the right solution to each problem instead of stretching the limits of a single approach while trying to force a match between this approach and a “non-fitting” problem (like using a screwdriver to knock a nail). Resulting from these considerations, multiple distinct data structures have been developed to tackle knowledge representation and exchange from various perspectives:

- Propositions;
- Graphs;
- Situation models;
- Spatial features;
- Hypotheses;
- Ontologies;
- Inference rules;
- Cases;
- Case templates;
- Kinematics and Geospatial Analysis Reasoning (KiGAR) configurations;
- Temporal Analysis and Reasoning (TAR) configurations; and
- Text-based templates.

Regarding knowledge representation requirements, the data structures listed above enable the representation of both situation knowledge and domain expert know-how (i.e., expertise). The former concerns knowledge about the various elements that constitute a situation of interest (e.g., persons, vehicles, buildings, organisations, weapons, etc.), the properties of these elements (e.g., age, colour, size, model, identification, etc.), and the relationships between the elements (e.g., is married with, works at, attacks, etc.). An example of domain expert know-how would be the anomaly detection knowledge of an analyst (typically acquired over years of experience) expressed as inference rules and/or cases previously encountered.

The data structures listed above also enable data, information and knowledge exchange between services (e.g., the WISDOM component themselves), agents and/or systems. When dealing with the exchange requirements, there are two main aspects to consider: 1) The actual "content" that has to be exchanged, and 2) The "container" that will be used to achieve the actual transfer of the content between components. The WISDOM data structures address the second aspect. In this regard, they have been designed to be generic and domain/source agnostic. That is, the content that can be exchanged is not prescribed by the data structures; it can be anything, from any domain. This really provides great flexibility for exploitation in various contexts and settings. From the data, information and knowledge exchange perspective, the instances of the data structures are documents in Extensible Markup Language (XML) format that constitute Data Transfer Objects (DTOs).

The WISDOM data structures can handle formal semantic data in addition to the typical data types (i.e., text, number, double, date, geometry, etc.). Hence, only when required, ontological formalism is used to handle semantics. Exploiting ontological formalism, when appropriate to do so, is considered a good practice, especially to support consistency across the system(s). However, it is not mandatory for the exploitation of the WISDOM data structures; it is optional. Moreover, one is not limited to a single ontology. Data from multiple ontologies (developed by and obtained from multiple providers) can be used simultaneously in the same instance of a WISDOM data structure. This also contributes to the exploitation flexibility.

Another convenient characteristic is that the data structures have been conceived to be compatible one with the other, as they exploit the same set of data types. Also of particular interest is that data conversion services have been developed to automatically convert a set of instances of the proposition data structure into a single instance of the graph data structure, and the opposite, i.e., convert a single graph into a set of propositions.

For usefulness, WISDOM provides the capability to create, retrieve, update and delete sets of instances of the proposition, spatial features, inference rule, case template and text-based template data structures.

Finally, regarding the data, information and knowledge persistence/storage requirements, a distributed approach exploiting heterogeneous database technologies has been adopted for WISDOM. A set of distinct data access services has been developed for this purpose. These services implement Create, Retrieve, Update and Delete (CRUD) operations on instances of the WISDOM data structures. Dedicated GUI components exist in WISDOM for the end-users to interact with these services.

D.2.1.1 Automated Reasoning Capability of the WISDOM R&D Platform

Automated inferencing/reasoning is a topic that has been devoted significant R&D efforts in a number of DRDC projects. All of these efforts have converged into the WISDOM R&D platform under the JICAC project. Figure D-1 illustrates the components of the current automated reasoning capability of WISDOM.

The capability is made of five reasoning services, each one implementing a distinct reasoning paradigm (rule-based, case-based, description logic, etc.). The multi-reasoner inferencing service can also be utilised as an orchestration service for the exploitation of the individual reasoning services. As can clearly be seen on Figure D-1, the proposition data structure plays a key role for data exchange in this automated reasoning capability; each automated reasoning service consumes propositions at its input, and it generates propositions (as results of the inferencing process) at its output.

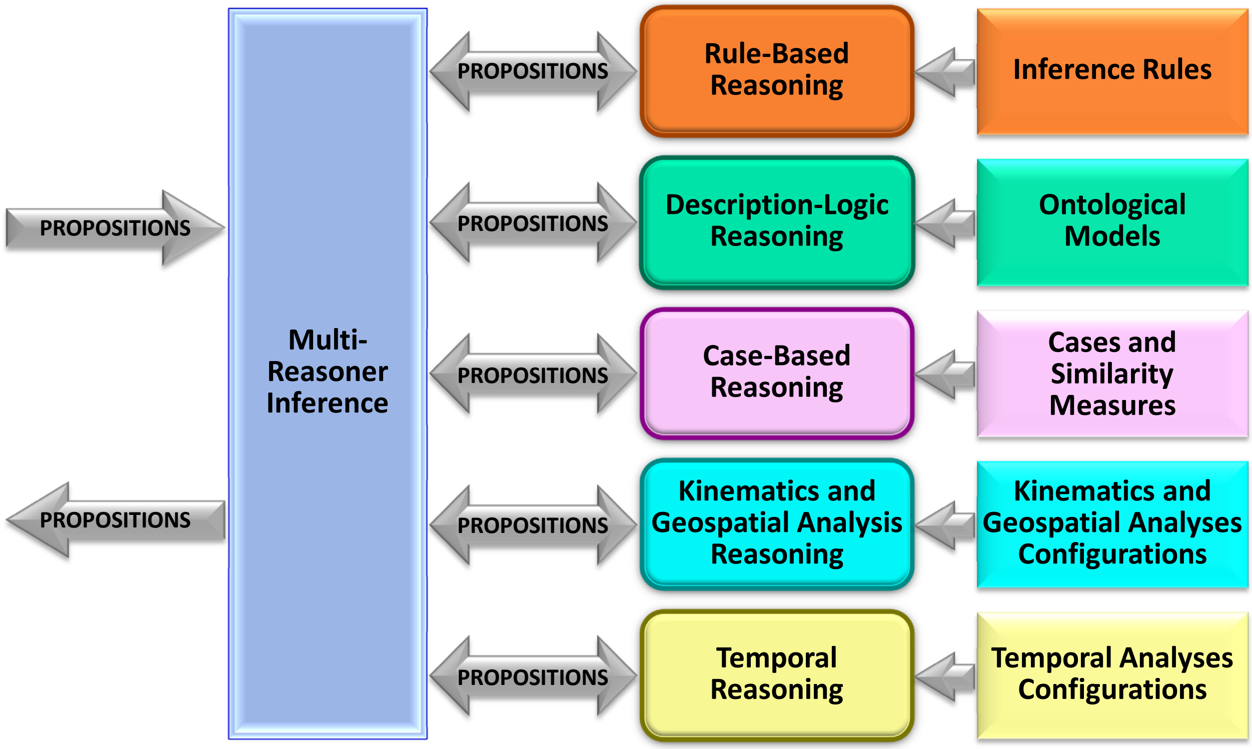

Figure D-1: WISDOM Automated Reasoning Capability.

D.3.1 REFERENCES

[1] Roy, J. (2020). WISDOM – A Platform for Proof-of-Concept R&D in Data and Information Integration, Fusion and Analytics. Defence Research and Development Canada (DRDC).

Annex E – UNCERTAINTY MANAGEMENT

In this annex, we provide a high level overview of the fundamentals of uncertainty management, in the context of KRR. This summary draws on a number of core references [1], [2], [3], [4].

Most tasks requiring intelligent behaviour have some degree of uncertainty associated with them. The uncertainty that can occur in knowledge-based systems may be caused by problems with the data. For example:

- Data might be missing or unavailable;
- Data might be present but unreliable or ambiguous due to measurement errors;
- The representation of the data may be imprecise or inconsistent;
- Data may just be user's best guess; and/or
- Data may be based on defaults and the defaults may have exceptions.

The uncertainty may also be caused by the represented knowledge, since it might:

- Represent best guesses of the experts that are based on plausible or statistical associations they have observed; and/or
- Not be appropriate in all situations (e.g., may have indeterminate applicability).

Given such numerous sources of errors, most knowledge-based systems require the incorporation of some form of uncertainty management.

When implementing an uncertainty scheme, one must be concerned with three issues:

- How to represent uncertain data;
- How to combine two or more pieces of uncertain data; and
- How to draw inference using uncertain data.

Through a review of the main typologies proposed in the literature, Jousselme et al. [1] consider a number of problems regarding the different types of uncertainty, the different epistemic interpretations, the different mathematical representations, in order to better understand and use the existing mathematical formalisms for reasoning under uncertainty.

E.1.1 UNCERTAINTY TYPOLOGY/TAXONOMY

Although it is possible to use semantic markup languages such as OWL to represent qualitative and quantitative information about uncertainty, there is no established foundation for doing so. Therefore, each developer must come up with his/her own set of constructs for representing uncertainty. This is a major liability in many environments that are dependent on interoperability among systems and applications. Moreover, apart from the interoperability issues caused by proprietary uncertainty representations, there are ancillary issues such as how to balance representational power vs. simplicity of uncertainty representations, which uncertainty representation technique(s) addresses use cases, how to ensure the consistency of representational formalisms and ontologies, etc.

Jousselme et al. [1] provide a discussion on uncertainty in the context of situation analysis. An overview of the principal typologies of uncertainty found in the literature at that time is presented, and the authors try to highlight useful distinctions. The wide array of uncertainty conceptions presented is a consequence of the intrinsic richness and ambiguity of natural language, but also a consequence of the complex nature of information.

Definitions of a limited number of concepts are provided in that work to illustrate and discuss the different facets of uncertainty. The benefits sought are:

- The avoidance of untimely uses of definitions and models of uncertainty;
- Clarifications allowing links with the already well developed logics of knowledge and belief; and
- Guidelines for the selection of the appropriate mathematical model to process uncertainty-based information.

E.2.1 WHAT IS UNCERTAINTY?

Uncertainty is a widely used term in the AI and engineering communities [5]. However, the authors in these fields of application and research do not always agree on the meaning of uncertainty, on its different types, on the possible sources, on the synonyms, on possible classifications, on representations, etc. In Laskey et al. [2] the term "uncertainty" is intended to encompass a variety of aspects of imperfect knowledge.

Jousselme et al. [1] attempt to clarify the concept of uncertainty and related concepts. They start with general definitions of uncertainty, provide a sociological point of view, and describe different visions of uncertainty that were proposed by some authors working in the AI and engineering areas.

Uncertainty has two main manifestations in most of the classical dictionaries.

- Uncertainty as a state of mind; and
- Uncertainty as a physical property of information.

The first refers to the state of mind of an agent, which does not possess the needed information or knowledge to make a decision; the agent is in a state of uncertainty: *"I'm not sure that this object is a table"*. The second refers to a physical property, representing the limitation of perception systems: *"The length of this table is uncertain"*.

In theories of uncertain reasoning, uncertainty is often described as imperfection of information, as errors on measures for example, and does not depend on any kind of state of mind. However, uncertain information can induce some uncertainty in our mind.

E.3.1 FORMALISMS FOR UNCERTAINTY MANAGEMENT

The term "*uncertainty reasoning*" is meant to denote the full range of methods designed for representing and reasoning with knowledge when Boolean truth values are unknown, unknowable, or inapplicable. To illustrate, consider a few reasoning challenges that could be addressed by reasoning under uncertainty:

- Automated agents are used to exchange information that in many cases is not perfect. Thus, a standardised format for representing uncertainty would allow agents receiving imperfect information to interpret it in the same way as was intended by the sending agents.
- Much information is likely to be uncertain. Examples include weather forecasts or gambling odds. Canonical methods for representing and integrating such information are necessary for communicating it in a seamless fashion.
- Information is also often incorrect or only partially correct, raising issues related to trust or credibility. Uncertainty representation and reasoning helps to resolve tension amongst information sources having different confidence and trust levels.

- Many visions rely on numerous distinct but conceptually overlapping ontologies that co-exist and interoperate. It is likely that in such scenarios, ontology mapping will benefit from the ability to represent degrees of membership and/or likelihoods of membership in categories of a target ontology, given information about class membership in the source ontology.
- Dynamic composability of services requires runtime identification of processing and data resources and resolution of policy objectives. Uncertainty reasoning techniques may be necessary to resolve situations in which existing information is not definitive.
- Information extracted from large information networks is typically incomplete. The ability to exploit partial information is very useful for identifying sources of service or information. It is clear that search effectiveness could be improved by appropriate use of technologies for handling uncertainty.

As work with semantics and services grows more ambitious, there is increasing appreciation of the need for principled approaches to representing and reasoning under uncertainty.

To model uncertainty, many mathematical tools have been developed, being either qualitative such as modal or nonmonotonic logics, or quantitative approaches such as probability theory, fuzzy sets theory, or evidential theory. These approaches are often compared on the basis of their different strengths and weaknesses:

- Their better suitability to model a particular type of uncertainty;
- Their requirement for prior knowledge;
- Their computational time complexity;
- The need for independence constraints; and
- Their reasoning capacities.

E.4.1 REFERENCES

[1] Jousselme, A.L., Maupin, P. and Bosse, E. (2003). Uncertainty in a Situation Analysis Perspective, IEEE. 2, pp. 1207-1214.

[2] Laskey, K., Laskey, K., Costa, P., Kokar, M., Martin T., and Lukasiewicz, T. (2008). Uncertainty Reasoning for the World Wide Web: Report on the URW3-XG Incubator Group. United States, North America.

[3] Costa, P.C.G., Laskey, K.B., Blasch, E. and Jousselme, A.L. (2012). Towards Unbiased Evaluation of Uncertainty Reasoning: The URREF Ontology, IEEE, pp. 2301-2308.

[4] University of Illinois. (2020). Chapter 4. Reasoning Under Uncertainty. Retrieved 01 May 2020, 2020, from https://www.cs.uic.edu/~liub/teach/cs511-spring-06/cs511-uncertainty.doc.

[5] Bosse, E., Roy, J. and Wark, S. (2007). Concepts, Models, and Tools for Information Fusion, Artech House.

Annex F – BIOGRAPHIES

Dr. Claire Bonial is a computational linguist specialising in the murky world of event semantics. In her efforts to make this world computationally tractable, she has collaborated on a variety of Natural Language Processing semantic role labelling projects, including PropBank, VerbNet, and Abstract Meaning Representation. A focussed contribution to these projects has been her theoretical and psycholinguistic research on both the syntax and semantics of English light verb constructions (e.g., *take a walk, make a mistake*). Bonial received her Ph.D. in Linguistics and Cognitive Science in 2014 from the University of Colorado Boulder. She began her current position in the Computational and Information Sciences Directorate of the Army Research Laboratory (ARL) in 2015. Since joining ARL, she has expanded her research portfolio to include multi-modal representations of events (text and imagery/video), as well as human-robot dialogue.

Jean Roy joined DRDC as a defence scientist in 1987. Since then, he has devoted his career to the automated integration, fusion, management and exploitation of data, information and knowledge, applied to situation and threat analysis and evaluation, in the context of command and control and intelligence, and also public security. For the last six years, he has been the lead scientist for the Joint Intelligence Collection and Analysis Capability project at DRDC, and also for the development of the WISDOM R&D platform. He has been a member of many international groups, panels and technical committees for various conferences, and he's received a TTCP achievement award. He is the author or co-author of over 150 publications, including the book "Concepts, Models and Tools for Information Fusion".

Dr. Forrest Hare is a retired Colonel in the United States Air Force most recently assigned to the Defense Intelligence Agency as the Deputy for the Indo-Asia Pacific Regional Center. As a major, he commanded an information warfare detachment in the European Air Operations Center during Operation Enduring Freedom. While assigned to the Air Staff Operations Directorate in the Pentagon, Dr. Hare was chosen to be on the Chief's Cyberspace Task Force to develop the vision for the Service's operations in its newest warfighting domain. His work contributed to the stand-up of the 24th Air Force and the creation of new cyberspace doctrine. After this assignment, he served on the staff of the Office of the Secretary of Defense and was a drafter of the Department of Defense Cyber Security policy. Dr. Hare has also served at the National Security Agency and in numerous overseas postings and deployments. In his current role as a solution architect for SAIC, he is developing a knowledge model for defence intelligence to improve the integration of cyber threat intelligence with traditional intelligence information.

Dr. Marielle Mokhtari gained a BSc in Computer Science in 1991 from Université Laval (Quebec City, Quebec, Canada), and a MSc and a PhD in Electrical Engineering with a specialisation in Computer Vision (data analytics), respectively in 1994 and 2000, both from Université Laval (Quebec City, Quebec, Canada). She worked for 2 years for the ABB Company as a senior system engineer, in the measurement and analytics sector, mainly in mathematical modelling and simulation of physical phenomena leading to the understanding of infrared spectrometers dedicated to the atmospheric research and to the military applications. She worked also in target detection in infrared images. Dr Mokhtari began as Defence Scientist for Defence Research and Development Canada Valcartier Research Centre (Quebec City, Quebec, Canada) in February 2002. Since that date, she continued to develop her knowledge and expertise in data analytics, visual analytics, information visualisation and scientific visualisation. She has also acquired knowledge of emerging technologies, namely virtual reality and augmented reality. She applies her expertise and knowledge to decision support related to various research areas of Defence and Security including: situational awareness in the context of contested urban environments; management and exploitation of information in tactical deployment; social network analysis for the military Intelligence; and understanding of a complex (tactical and operational) situation.

Dr. Hans-Christian Schmitz, Fraunhofer FKIE (Germany): Hans-Christian earned his PhD in computational linguistics at the University of Bonn and worked for various universities and research organisations with a main focus on natural language processing, knowledge representation and data science. He is now a researcher at Fraunhofer FKIE, where he is involved in various projects within the areas of semantic interoperability and data analytics. In particular, he is leading a project team that investigates the application of AI-methodologies – including both ML and symbolic reasoning – for C2 support.

Mr. John Sweet has over 30 years of experience in supporting geospatial and all-source intelligence as an analyst, engineer and architect. His work history includes Soviet and post-Soviet analysis with USEUCOM, management of the Linked Operations-Intelligence Center Europe (LOCE) Southern European Command environment during the crisis in Bosnia and Kosovo, Imagery Analysis training for US National Geospatial-Intelligence Agency (NGA), Ground Station Operations for the US National Reconnaissance Organization (NRO), and coalition partner interoperability/standards for the US Office of the Undersecretary of Defense for Intelligence (OUSD-I). Mr Sweet is currently acting in the role of Partner Integrator to the NGA Structured Observation Management (SOM) program.

Dr. David Barber is currently the principal technical authority for the Underpinning Data Science project at the UK’s Defence Science and Technology Laboratory (Dstl). Since joining Dstl he has held roles leading strategy for the Intelligence Systems and Analytics Group, as a team leader and as a technical consultant supporting MOD’s geospatial enterprise. Through those roles he has led on academic engagement with UK partners such as the Alan Turing Institute, support to Joint Forces Command and Defence Equipment and Support. David’s technical expertise is in geospatial data management, interoperability and exploitation with interests in data science and AI. More broadly he enjoys building collaboration across diverse groups and understanding how technology can make a real difference to customer’s challenges. Prior to joining Dstl David was a postdoctoral research associate at Newcastle University, from where he also holds a PhD. He is a Chartered Engineer.

Dr. Paul Cripps specialises in information and knowledge organisation systems, particularly those of a geospatial or spatio-temporal nature. His PhD focussed on the application of semantic technologies for research purposes. Current research includes the development and application of phenomenological, observational approaches to spatio-temporal knowledge graphs as well as the application of knowledge organisation systems for knowledge management and the development and application of geospatial information systems and standards.

Dr. Chris Mowat is a work package leader for the Digital Data Deception project at Dstl, as well as a data scientist for the Underpinning Data Science project. Chris’s research focuses on aspects of AI and data science including the vulnerability of AI systems to deliberate deception, and knowledge representation to support information advantage. Prior to joining Dstl in 2019, Chris obtained a PhD in astronomy from the University of Exeter.

REPORT DOCUMENTATION PAGE

1. Recipient's Reference	2. Originator's References	3. Further Reference	4. Security Classification of Document
	STO-TR-IST-ET-111 AC/323(IST-111)TP/1022	ISBN 978-92-837-2342-4	PUBLIC RELEASE

5. Originator Science and Technology Organization
North Atlantic Treaty Organization
BP 25, F-92201 Neuilly-sur-Seine Cedex, France

6. Title

Knowledge Representation and Reasoning – A Review of the State of the Art and Future Opportunities

7. Presented at/Sponsored by

A report from NATO IST ET-111.

8. Author(s)/Editor(s)	9. Date
Multiple	June 2022

10. Author's/Editor's Address	11. Pages
Multiple	98

12. Distribution Statement There are no restrictions on the distribution of this document. Information about the availability of this and other STO unclassified publications is given on the back cover.

13. Keywords/Descriptors

Big data; Causality; Explainability; Exploitation; Knowledge engineering; Knowledge representation; Ontology; Reasoning; Sense-making; Trust

14. Abstract

This report presents the findings of NATO Information Systems Technology (IST) Exploratory Team 111 (ET-111) regarding the status of Knowledge Representation and Reasoning (KRR). Knowledge representation is the expression of knowledge that is computer-tractable, able to be reasoned over, and exploitable – especially in the age of big data.

This report provides a technical introduction to the field of KRR. The capabilities of knowledge systems are discussed, as well as how they can be created, and how they may then be applied to real-world military problems. The distinction is made between KRR ('symbolic' AI) and machine learning ('sub-symbolic' AI), along with a consideration of how the two approaches can be complementary.

We discuss some of the challenges for NATO member nations, and how KRR can influence these areas. We present specific examples of knowledge representation including the Multilateral Interoperability Program (MIP) Information Model (MIM), the Rich Event Ontology (REO) and the Defense Intelligence Core Ontology (DICO). Moreover, we review areas of research including causality, explainability, and trust in KRR systems.

Finally, we discuss implications for NATO, and present our key recommendations for future work involving technical activities and exploratory teams.

NORTH ATLANTIC TREATY ORGANIZATION

SCIENCE AND TECHNOLOGY ORGANIZATION

BP 25
F-92201 NEUILLY-SUR-SEINE CEDEX • FRANCE
Télécopie 0(1)55.61.22.99 • E-mail mailbox@cso.nato.int

DIFFUSION DES PUBLICATIONS
STO NON CLASSIFIEES

Les publications de l'AGARD, de la RTO et de la STO peuvent parfois être obtenues auprès des centres nationaux de distribution indiqués ci-dessous. Si vous souhaitez recevoir toutes les publications de la STO, ou simplement celles qui concernent certains Panels, vous pouvez demander d'être inclus soit à titre personnel, soit au nom de votre organisation, sur la liste d'envoi.

Les publications de la STO, de la RTO et de l'AGARD sont également en vente auprès des agences de vente indiquées ci-dessous.

Les demandes de documents STO, RTO ou AGARD doivent comporter la dénomination « STO », « RTO » ou « AGARD » selon le cas, suivi du numéro de série. Des informations analogues, telles que le titre est la date de publication sont souhaitables.

Si vous souhaitez recevoir une notification électronique de la disponibilité des rapports de la STO au fur et à mesure de leur publication, vous pouvez consulter notre site Web (http://www.sto.nato.int/) et vous abonner à ce service.

CENTRES DE DIFFUSION NATIONAUX

ALLEMAGNE
Streitkräfteamt / Abteilung III
Fachinformationszentrum der Bundeswehr (FIZBw)
Gorch-Fock-Straße 7, D-53229 Bonn

BELGIQUE
Royal High Institute for Defence – KHID/IRSD/RHID
Management of Scientific & Technological Research for Defence, National STO Coordinator
Royal Military Academy – Campus Renaissance
Renaissancelaan 30, 1000 Bruxelles

BULGARIE
Ministry of Defence
Defence Institute "Prof. Tsvetan Lazarov"
"Tsvetan Lazarov" bul no.2
1592 Sofia

CANADA
DGSIST 2
Recherche et développement pour la défense Canada
60 Moodie Drive (7N-1-F20)
Ottawa, Ontario K1A 0K2

DANEMARK
Danish Acquisition and Logistics Organization (DALO)
Lautrupbjerg 1-5
2750 Ballerup

ESPAGNE
Área de Cooperación Internacional en I+D
SDGPLATIN (DGAM)
C/ Arturo Soria 289
28033 Madrid

ESTONIE
Estonian National Defence College
Centre for Applied Research
Riia str 12
Tartu 51013

ETATS-UNIS
Defense Technical Information Center
8725 John J. Kingman Road
Fort Belvoir, VA 22060-6218

FRANCE
O.N.E.R.A. (ISP)
29, Avenue de la Division Leclerc
BP 72
92322 Châtillon Cedex

GRECE (Correspondant)
Defence Industry & Research General Directorate, Research Directorate
Fakinos Base Camp, S.T.G. 1020
Holargos, Athens

HONGRIE
Hungarian Ministry of Defence
Development and Logistics Agency
P.O.B. 25
H-1885 Budapest

ITALIE
Ten Col Renato NARO
Capo servizio Gestione della Conoscenza
F. Baracca Military Airport "Comparto A"
Via di Centocelle, 301
00175, Rome

LUXEMBOURG
Voir Belgique

NORVEGE
Norwegian Defence Research Establishment
Attn: Biblioteket
P.O. Box 25
NO-2007 Kjeller

PAYS-BAS
Royal Netherlands Military Academy Library
P.O. Box 90.002
4800 PA Breda

POLOGNE
Centralna Biblioteka Wojskowa
ul. Ostrobramska 109
04-041 Warszawa

PORTUGAL
Estado Maior da Força Aérea
SDFA – Centro de Documentação
Alfragide
P-2720 Amadora

REPUBLIQUE TCHEQUE
Vojenský technický ústav s.p.
CZ Distribution Information Centre
Mladoboleslavská 944
PO Box 18
197 06 Praha 9

ROUMANIE
Romanian National Distribution Centre
Armaments Department
9-11, Drumul Taberei Street
Sector 6
061353 Bucharest

ROYAUME-UNI
Dstl Records Centre
Rm G02, ISAT F, Building 5
Dstl Porton Down
Salisbury SP4 0JQ

SLOVAQUIE
Akadémia ozbrojených síl gen. M.R. Štefánika, Distribučné a informačné stredisko STO
Demänová 393
031 01 Liptovský Mikuláš 1

SLOVENIE
Ministry of Defence
Central Registry for EU & NATO
Vojkova 55
1000 Ljubljana

TURQUIE
Milli Savunma Bakanlığı (MSB)
ARGE ve Teknoloji Dairesi Başkanlığı
06650 Bakanliklar – Ankara

AGENCES DE VENTE

The British Library Document Supply Centre
Boston Spa, Wetherby
West Yorkshire LS23 7BQ
ROYAUME-UNI

Canada Institute for Scientific and Technical Information (CISTI)
National Research Council Acquisitions
Montreal Road, Building M-55
Ottawa, Ontario K1A 0S2
CANADA

Les demandes de documents STO, RTO ou AGARD doivent comporter la dénomination « STO », « RTO » ou « AGARD » selon le cas, suivie du numéro de série (par exemple AGARD-AG-315). Des informations analogues, telles que le titre et la date de publication sont souhaitables. Des références bibliographiques complètes ainsi que des résumés des publications STO, RTO et AGARD figurent dans le « NTIS Publications Database » (http://www.ntis.gov).

NORTH ATLANTIC TREATY ORGANIZATION

NATO OTAN

BP 25
F-92201 NEUILLY-SUR-SEINE CEDEX • FRANCE
Télécopie 0(1)55.61.22.99 • E-mail mailbox@cso.nato.int

SCIENCE AND TECHNOLOGY ORGANIZATION

DISTRIBUTION OF UNCLASSIFIED STO PUBLICATIONS

AGARD, RTO & STO publications are sometimes available from the National Distribution Centres listed below. If you wish to receive all STO reports, or just those relating to one or more specific STO Panels, they may be willing to include you (or your Organisation) in their distribution.

STO, RTO and AGARD reports may also be purchased from the Sales Agencies listed below.

Requests for STO, RTO or AGARD documents should include the word 'STO', 'RTO' or 'AGARD', as appropriate, followed by the serial number. Collateral information such as title and publication date is desirable.

If you wish to receive electronic notification of STO reports as they are published, please visit our website (http://www.sto.nato.int/) from where you can register for this service.

NATIONAL DISTRIBUTION CENTRES

BELGIUM
Royal High Institute for Defence – KHID/IRSD/RHID
Management of Scientific & Technological Research for Defence, National STO Coordinator
Royal Military Academy – Campus Renaissance
Renaissancelaan 30
1000 Brussels

BULGARIA
Ministry of Defence
Defence Institute "Prof. Tsvetan Lazarov"
"Tsvetan Lazarov" bul no.2
1592 Sofia

CANADA
DSTKIM 2
Defence Research and Development Canada
60 Moodie Drive (7N-1-F20)
Ottawa, Ontario K1A 0K2

CZECH REPUBLIC
Vojenský technický ústav s.p.
CZ Distribution Information Centre
Mladoboleslavská 944
PO Box 18
197 06 Praha 9

DENMARK
Danish Acquisition and Logistics Organization (DALO)
Lautrupbjerg 1-5
2750 Ballerup

ESTONIA
Estonian National Defence College
Centre for Applied Research
Riia str 12
Tartu 51013

FRANCE
O.N.E.R.A. (ISP)
29, Avenue de la Division Leclerc – BP 72
92322 Châtillon Cedex

GERMANY
Streitkräfteamt / Abteilung III
Fachinformationszentrum der Bundeswehr (FIZBw)
Gorch-Fock-Straße 7
D-53229 Bonn

GREECE (Point of Contact)
Defence Industry & Research General Directorate, Research Directorate
Fakinos Base Camp, S.T.G. 1020
Holargos, Athens

HUNGARY
Hungarian Ministry of Defence
Development and Logistics Agency
P.O.B. 25
H-1885 Budapest

ITALY
Ten Col Renato NARO
Capo servizio Gestione della Conoscenza
F. Baracca Military Airport "Comparto A"
Via di Centocelle, 301
00175, Rome

LUXEMBOURG
See Belgium

NETHERLANDS
Royal Netherlands Military Academy Library
P.O. Box 90.002
4800 PA Breda

NORWAY
Norwegian Defence Research Establishment, Attn: Biblioteket
P.O. Box 25
NO-2007 Kjeller

POLAND
Centralna Biblioteka Wojskowa
ul. Ostrobramska 109
04-041 Warszawa

PORTUGAL
Estado Maior da Força Aérea
SDFA – Centro de Documentação
Alfragide
P-2720 Amadora

ROMANIA
Romanian National Distribution Centre
Armaments Department
9-11, Drumul Taberei Street
Sector 6
061353 Bucharest

SLOVAKIA
Akadémia ozbrojených síl gen M.R. Štefánika, Distribučné a informačné stredisko STO
Demänová 393
031 01 Liptovský Mikuláš 1

SLOVENIA
Ministry of Defence
Central Registry for EU & NATO
Vojkova 55
1000 Ljubljana

SPAIN
Área de Cooperación Internacional en I+D
SDGPLATIN (DGAM)
C/ Arturo Soria 289
28033 Madrid

TURKEY
Milli Savunma Bakanlığı (MSB)
ARGE ve Teknoloji Dairesi Başkanlığı
06650 Bakanliklar – Ankara

UNITED KINGDOM
Dstl Records Centre
Rm G02, ISAT F, Building 5
Dstl Porton Down, Salisbury SP4 0JQ

UNITED STATES
Defense Technical Information Center
8725 John J. Kingman Road
Fort Belvoir, VA 22060-6218

SALES AGENCIES

The British Library Document Supply Centre
Boston Spa, Wetherby
West Yorkshire LS23 7BQ
UNITED KINGDOM

Canada Institute for Scientific and Technical Information (CISTI)
National Research Council Acquisitions
Montreal Road, Building M-55
Ottawa, Ontario K1A 0S2
CANADA

Requests for STO, RTO or AGARD documents should include the word 'STO', 'RTO' or 'AGARD', as appropriate, followed by the serial number (for example AGARD-AG-315). Collateral information such as title and publication date is desirable. Full bibliographical references and abstracts of STO, RTO and AGARD publications are given in "NTIS Publications Database" (http://www.ntis.gov).

ISBN 978-92-837-2342-4