

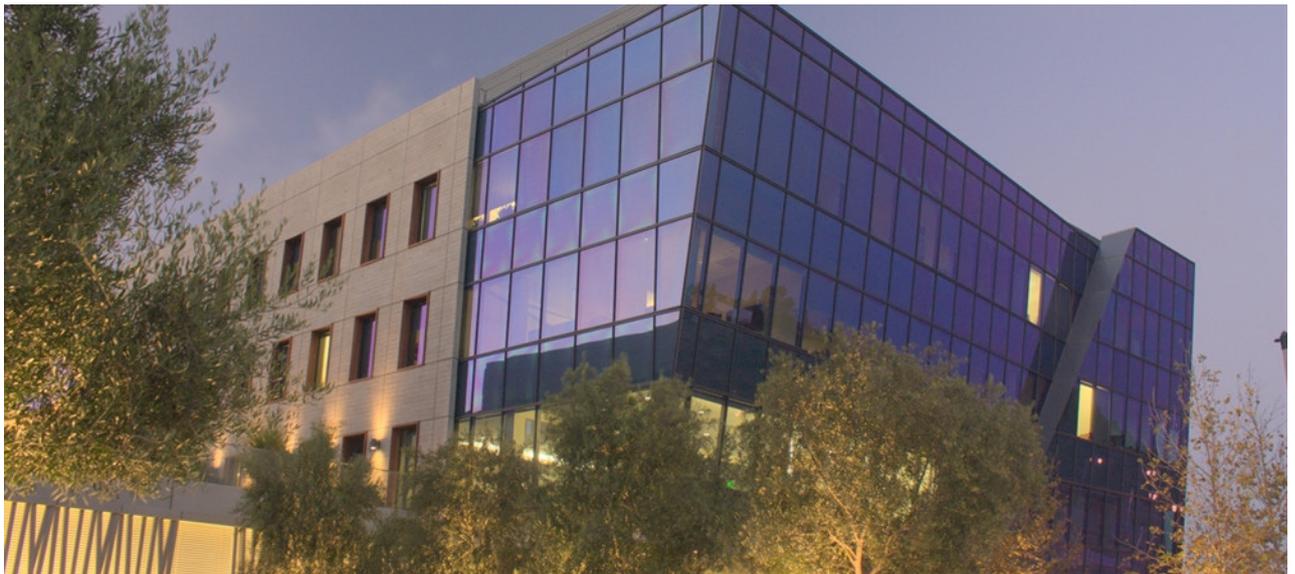
KSCO

Proceedings of the Ninth International Conference on Knowledge Systems for Coalition Operations 2017

6th to 8th November 2017 in Los Angeles, California, USA

Co-located with the 22nd International Command and Control Research and
Technology Symposium (ICCRTS 2017)

Army Research Laboratory West and USC Institute for Creative Technologies
12015 Waterfront Drive, Playa Vista, CA 90094-2536, USA



These proceedings are also available via the KSCO-2017 and ICCRTS-2017 web sites

<http://ksco.info/ksco-2017.html>
<http://internationalc2institute.org/proceedings/>

Editor: Austin Tate, University of Edinburgh

© 2017, Copyright remains with the authors of the individual papers.



Knowledge Systems for Coalition Operations 2017 Proceedings

6th to 8th November 2017 in Los Angeles, California, USA

The Ninth Knowledge Systems for Coalition Operations Conference (KSCO-2017) was incorporated into the 22nd International Command and Control Research and Technology Symposium (ICCRTS 2017)

Introduction

KSCO - Knowledge Systems for Coalition Operations is an international working group exploring research in Knowledge-based Systems and Information Management, with a focus on the challenges of Coalition Operations. KSCO regularly organizes a technical conference where practitioners and key decision makers in coalition operations management meet and discuss with researchers from areas of knowledge-based systems, information management, planning, and multi-agent systems, exchange experience and ideas, share inspiration, and suggest novel concepts. It can also lead to joint project proposals.

After successful events in Edinburgh UK, Toulouse France, Prague Czech Republic, Waltham MA USA, Southampton UK, Vancouver BC Canada, Pensacola Florida USA, and London UK, the ninth KSCO conference in 2017 was held in Los Angeles California USA in conjunction with the 22nd International Command and Control Research and Technology Symposium (ICCRTS).

Topics of Interest

KSCO-2017 is a forum to publish original research, application and project description papers related to intelligent and knowledge systems for coalition operations management. A coalition includes, but not limited to, military, inter-agency and cross organization alliances engaged in a cooperative endeavour and joining capabilities together for a common cause. Topics may be related to knowledge systems requirements and knowledge systems potential or actual use for coalition and inter-agency operations (i.e., multi-national and multi-agency, civil authorities, home land safety and security, expeditionary or domestic operations).

KSCO 2017 areas include:

- Intelligent Command and Control (C2)
- Coalition and Team Information Sharing
- Coalition and Team User Interfaces
- Planning and Scheduling
- Coordination and Collaboration
- Data to Decision Support

- Humanitarian Assistance & Disaster Recovery
- Cultural Influences
- Human-Machine Teams
- Virtual Coalition Organizations
- KSCO related Research Programmes & Projects
- Deployed Systems & Case Studies

Organizing Committee

- ICCRTS-2017 Conference Chair: Dr. David Alberts, IC2I, USA
- ICCRTS-2017 Programme Chair: Dr. Micheline Bélanger, DRDC, Canada
- ICCRTS-2017/KSCO-2017 Local Arrangements Chair: Dr. Niranjan Suri, IHMC, Florida, USA
- KSCO-2017 Conference Chair: Dr. Jitu Patel, Dstl, UK
- KSCO-2017 Programme Chair & Publications Chair: Prof. Austin Tate, AIAI, Edinburgh, UK
- KSCO-2017 Organizing Committee Members:
 - Dr. Liz Bowman, ARL, USA
 - Dr. Tim Grant, R-BAR, Netherlands
 - Mr. Rick Metzger, AFRL, USA
 - Prof. Michal Pěchouček, Czech Technical University, Prague, Czech Republic

3: Austin Tate (UK, University of Edinburgh) and Jeffrey Hansberger (USA, Army Research Laboratory)

Virtual Operations Centres for Coalition Operations and Distributed Team Collaboration

Abstract:

On-line multi-user virtual worlds have been used to create collaboration environments and shared virtual spaces to allow distributed teams to train, exercise or work together. Platforms such as Linden Lab's Second Life or the open source community's OpenSimulator have been used to provide easily accessed facilities in which users are represented by avatars in a space designed to support their collaboration and sharing of resources. The creation of a suitable virtual space allows users wherever they are located to be brought together into a shared visualisation of an "operations centre". This may be joined with real operations centre(s) to integrate a distributed team to allow them to more effectively address the task or operations they are engaged in. Such environments are particularly well suited to training and exercises, but can also be used for real events when distributed teams are involved.

The paper describes the "Open Virtual Collaborative Environment" (OpenVCE) and its facilities, and how the resources have been made widely available as a basis for creating customised environments and used for multi-national and multi-agency team collaboration facilities especially where teams are geographically distributed.

Download: [Paper PDF](#), [Presentation PDF](#)

72: Wen Zhou (China, National University of Defense Technology), Weidong Bao (China, National University of Defense Technology), Xiaomin Zhu (China, National University of Defense Technology), Chao Chen (China, National University of Defense Technology), Zhu Wu (China, Naval Command College) and Changfei Wu (Air Force Early Warning Academy)

A Semantic-Based Information Integration Framework of Agile Command and Control

Abstract:

Human society is in the period of accelerated development of the information revolution. It not only affects people's life extensively, but also constantly updates our understanding and understanding of the mode of operation in the information age. The ability to carry out effective information integration and provide intelligent information services, has become a joint combat process to obtain the command and control of the agile advantage of the support base. Based on this, this paper proposes a semantic-based information integration framework of agile command and control, which includes two parts: basic module and semantic analysis engine. The basic idea is to change the semantic relation from the pursuit of comprehensive information reasoning to the attention data, the idea of information retrieval into knowledge retrieval, relying on the information integration framework to tap the command information

system implied in a variety of rich semantic relations, in order to carry out a higher level of information integration and provide more intelligent knowledge retrieval, and thus agile Command and control.

76: Jitu Patel (UK, Dstl) and Peter Houghton (UK, Dstl)

Where are the Knowledge Systems? - Understanding Obstacles to Technology Adoption

Abstract:

Command and Control (C2) aspects of military operations, though very complex, knowledge intensive and usually time bounded, tend to be carried out with limited use of command and decision support tools. While there have been isolated successes such as the Dynamic Analysis and Replanning Tool used during operation Desert Shield, routine adoption of knowledge systems have remained low. The problem is not with the availability of technology, as defence (both Governments and Industries) have invested significant effort in tool development over the years. The Aim of this paper is to discuss findings from an earlier study, which investigated factors that may be contributing to the lack of exploitation of the intended utility of knowledge systems, and also to present insights from recent experience working in a standing joint force headquarters. The paper will conclude with an agenda for future research that contributes to our understanding of how best to improve technology adoption.

111: David Mott (UK, Emerging Technology Services, IBM UK), Yunfeng Zhang (USA, Human Agent Collaboration Group, IBM T.J. Watson Research Center), Soheil Eshghi (USA: Department of Electrical Engineering and Institute for Network Science, Yale University), Cheryl Giammanco (USA, Army Research Laboratory) and Troy Kelley (USA, Human Research and Engineering Directorate, Army Research Laboratory)

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

Abstract:

We describe research on understanding group mutability in the behaviour of external groups, and how interventions by coalition forces may affect the behaviour in terms of controlling hostile groups and encouraging friendly groups. We explore how emotion may influence the behaviour of individuals by affecting the type of reasoning that they undertake, encouraging "uncritical" rather than "critical" thinking. We describe a computational framework holding a cognitive model of an individual operating within a group context, inspired by theories from social science. Individuals relate to in-groups and out-groups and have beliefs that are associated with emotions. Cognitive Appraisal Theory is used to evaluate incoming memes "pronounced" by external speakers, appraising the effects of the memes on an individual's self-esteem taking account of their group relationships as indicated by social identity theory, and leading to an emotion in the individual. Appraisal is followed by a process of coping that seeks to handle the effects by either performing problem-focussed (critical) or emotion-focussed (uncritical) thinking, according to the current emotional state of the individual. This

model is implemented within a Cognitive Architecture (SOAR) as a set of reasoning processes that handle beliefs and emotion. The model is integrated into a multi-agent simulation tool (Repast Simphony) allowing the simulation of populations of individuals interacting and spreading rumours, or memes, together with interventions. We describe how this framework could be used to construct experiments to explore how different situations lead to group mutability and behaviour, together with the effects of interventions by coalition forces.

118: Karen Myers (USA, SRI International), Tim Ellis (USA, SRI International) and Tancrede Lepoint (USA, SRI International)

Privacy Technologies for Controlled Information Sharing in Coalition Operations

Abstract:

Information sharing among coalition partners must balance the benefits that can accrue from improved coordination with the risks of releasing information that ideally would be kept private. We consider how advanced privacy technologies can enable improved information sharing among coalition partners by both providing increased control over how information is used or released, and enabling principled characterizations of the impact of individual and cumulative sharing activities. We describe this work in the context of a humanitarian aid and disaster relief (HADR) scenario, showing how the technologies can enable significantly increased and informed sharing.

121: Erica P. Viklund (USA, Pacific Science & Engineering Group), Heather M. Oonk (USA, Pacific Science & Engineering Group) and Manuela Jaramillo (USA, Pacific Science & Engineering Group)

The Role of Transactive Memory (TM) in Proactive Decision Support (PDS)

Abstract:

Advances in technology have exponentially increased the information and data at our fingertips. While there are many benefits of such access, a tradeoff is that information seekers can be overwhelmed by the vast sea of information at their disposal. Challenges multiply when information seekers operate as part of a team where there are differences in knowledge, information access, and decision-making responsibilities. Coalition operations are examples of such situations, involving decisions that impact a complicated network of different countries and actors. Proactive decision support (PDS) tools have the potential to make more manageable the tasks of selecting, verifying, compiling, and analyzing relevant information, so that good decisions can be made more efficiently. Effective PDS requires a system that "understands" and adapts to the context in which information seeking and decision-making occur. Context includes aspects of the physical environment within which the technology and user are embedded, and the cognitive or mission objectives of users. We argue that for teams, PDS context must also include a collection of team member and team dynamic variables such as shared and differential tasks, requirements, knowledge, and expertise. Collectively, these

variables can be conceptualized as transactive memory (TM). We describe how PDS that incorporates TM variables as a form of context can facilitate and streamline validation and communication of information among team members, which is crucial for realizing the potential benefits of PDS for coalition operations. We discuss considerations for implementing TM variables into PDS tools and key research and development questions to be addressed.

122: Adrienne Raglin (USA, Army Research Lab) and Christian Schlesiger (USA, Army Research Lab)

Data Analyzer Software: a Knowledge System Supporting Coalition and Team Information Sharing

Abstract:

As current and future operations integrate soldiers from multiple nations, information that supports short term and long term teaming is critical. Among coalition forces it is important to maintain unity of effort, to plan concurrently, and to make adjustment in sync ensuring operations are carried out successfully. Combatant commanders have many responsibilities including ensuring the capability and capacity of the forces with partnering nations.

However, in multinational operations there is the added need to consider differences in organization, doctrine, terminology, and objectives. This can be achieved through knowledge capturing, information sharing, and training. Additionally, giving commanders required information with explanation, linking knowledge and uncertainty can improve teamed operations in complex and dynamic environments. The Data Analyzer was initially designed for data analysis of training software previously developed at the Army Research Laboratory. Now that this training software is being used by US and coalition partners, the Data Analyzer has been expanded as a platform for wider analysis and knowledge capturing.

The Data Analyzer provides commanders with the ability to view data capturing detailed experiential knowledge and find trends in tactics, techniques and procedures (TTPs) employed within their units and within different coalition partners. This information from the analyzer provides the joint forces with similarities and differences highlighted that aid in joint engagement preparation and insights into actions that can impact joint mission TTPs. We present the Data Analyzer software and use case scenarios illustrating utilizing this approach in supporting knowledge capturing, information sharing, and decision making for multinational operations.

Virtual Operations Centres for Coalition Operations and Distributed Team Collaboration

Austin Tate <a.tate@ed.ac.uk>

AIAI, University of Edinburgh, Informatics Forum, Crichton Street, Edinburgh EH8 9AB, U.K.

Jeffrey T. Hansberger <jeffrey.t.hansberger.civ@mail.mil>

U.S. Army Research Laboratory HRED, Redstone Arsenal, Alabama 35898-7290, U.S.A.

Abstract

On-line multi-user virtual worlds have been used to create collaboration environments and shared virtual spaces to allow distributed teams to train, exercise or work together. Platforms such as Linden Lab's Second Life™ or the open source community's OpenSimulator have been used to provide easily accessed facilities in which users are represented by avatars in a space designed to support their collaboration and sharing of resources. The creation of a suitable virtual space allows users wherever they are located to be brought together into a shared visualisation of an "operations centre". This may be joined with real operations centre(s) to integrate a distributed team to allow them to more effectively address the task or operations they are engaged in. Such environments are particularly well suited to training and exercises, but can also be used for real events when distributed teams are involved.

The paper describes the "Open Virtual Collaborative Environment" (OpenVCE) and its facilities, and how the resources have been made widely available as a basis for creating customised environments and used for multi-national and multi-agency team collaboration facilities especially where teams are geographically distributed.

Introduction

On-line multi-user virtual worlds have been used to create collaboration environments and shared virtual spaces to allow distributed teams to train, exercise or work together. Platforms such as Linden Lab's Second Life™ or the open source community's OpenSimulator have been used to provide easily accessed facilities in which users are represented by avatars in a space designed to support their collaboration and sharing of resources. The creation of a suitable virtual space allows users wherever they are located to be brought together into a shared visualisation of an “operations centre”. This may be joined with real operations centre(s) to integrate a distributed team to allow them to more effectively address the task or operations they are engaged in. Such environments are particularly well suited to training and exercises, but can also be used for real events when distributed teams are involved (Tate, 2006).

This paper describes work on the “Open Virtual Collaboration Environment” (OpenVCE), a project which explored openly accessible platforms, tools and protocols to support distributed team collaboration. The paper has a focus particularly on the provision of virtual world operations centres for both training exercises and actual emergency response.

Emergency Response

The desire to support the collaborative development of responses to large-scale emergency crises provided the impetus for the work described here, although the lessons learned should be applicable to other types of interaction. Crisis response situations require collaboration among individuals belonging to many different organizations and having different backgrounds, training, procedures and objectives. The response to the Indian Ocean Tsunami in 2004 and the Hurricane Katrina relief efforts in 2005 emphasized the importance of effective communication and collaboration. In the former, the Multinational Planning Augmentation Team (MPAT) supported brokering of requests for assistance by matching them with offers of help from deployed military and humanitarian assistance facilities. In the aftermath of Hurricane Katrina, the U.S. Army and National Guard assisted state, federal, and non-government organizations with varying degrees of efficiency and expediency. Compounding the challenges associated with such situations is the distributed nature of the community of experts who can contribute to the analysis of the crisis and the planning of a response. As a result, opportunities for leveraging expertise and resources across organizations are haphazard at best, and the response to the crisis can appear as chaotic as the crisis itself.

Seeking more effective and efficient means to facilitate crisis response, in 2009 the U.S. Joint Forces Command (USJFCOM) and the U.S. Army Research Laboratory's Human Research and Engineering Directorate (ARL HRED) launched a project under the direction of one of the authors (Hansberger) to design and evaluate a Virtual Collaboration Environment (VCE), and to seek to demonstrate its potential for distributed crisis response planning. More broadly, the project sought to discover implications for any distributed collaborative activity.

Virtual Collaboration Environment Concept

The initial technical concept behind the VCE was to investigate the potential of new media technologies, specifically social networking and virtual worlds, to provide a virtual environment that fosters community spirit and collaborative effort in some particular field (a field in which, we assume, there exists a potential community of users who have complementary knowledge or skills that contribute to problem-solving). Thus envisaged, the VCE was planned to meet several specific requirements:

- The creation and maintenance of a community of on-line users with diverse backgrounds (including those with little or no prior experience of virtual or on-line communities). In the first instance, the VCE was intended to support a "Whole of Society Crisis Response" (WoSCR) community, a loosely affiliated community of subject-matter experts and crisis responders drawn from international government and civilian organizations for the purpose of contributing their specialised knowledge to crisis response planning activities. In the course of the project an initial mailing list of 1,600 people already involved in such activities was used to establish the community, of which some 300 were active within the VCE facilities that were provided. It contained members from a number of countries (although initially with a strong U.S. bias) drawn from the worlds of government, business and academia.
- The ability for users to conduct *synchronous* collaborations for the purpose of collective decision-making during specific problem-solving episodes.
- The provision for the users of mechanisms for the *asynchronous* creation and development of on-line material. This has two aspects: the short-term development of informational material as a part of the problem-solving process; and the long-term development of an on-line body of experience, knowledge and debate about the field in question.

Hence, collaborations in the environment would have two, quite different, aspects: a continuous asynchronous collaboration among users to discuss and develop on-line documentation pertaining to their field of interest (activities which would also help foster a sense of community); and interspersed synchronous problem-solving collaborations of relatively short duration in which their expertise is put into practice. It was envisaged that web and virtual worlds' technologies together would provide the technical backbone for meeting these requirements.

Virtual Collaboration Environment Requirements and Design

The designers and developers of the VCE included groups from the University of Edinburgh, the University of Virginia, Carnegie Mellon University and Perigeon Technologies LLC, each of which had an existing and complementary interest in collaborative work and so would bring specialised knowledge or technology to the programme. Therefore, the U.S. Army Research Laboratory and these partners engaged in the creation of an openly accessible "Virtual Collaborative Environment" (VCE) to support the "Whole of Society Crisis Response" (WoSCR) community of interest involved in crisis action planning and execution activities (Hansberger et al., 2010) and also in the later work to support the "Dismounted Infantry

Collaboration Environment" (DICE) for remote support to medics supporting injured soldiers (Tate et al., 2012).

The VCE consists of a collaborative portal containing a suite of Web 2.0 social networking and group support tools including data visualisation facilities (Moon et al., 2011), a 3D virtual world collaboration space (Tate et al., 2010a) and a virtual collaboration protocol based on social science research (Cross and Parker, 2004) to assist the team members to work together effectively. The aim was to choose from and utilise appropriate capabilities from this suite of tools to meet the requirements of distributed collaboration for the target communities.

As a first step, in order to validate these initial assumptions, a Cognitive Work Analysis (CWA) (Vicente, 1999) was performed. All tools were selected to support the key functions identified in a Work Domain Analysis (figure 1) for distributed collaboration (more fully reported on in Tate et al., 2014). The CWA bottoms out in a number of types of tool or technology which “facilitate” the required communications methods and activities. This set of requirements and types of technology is used to guide selection and provision of key features in the experimental collaboration environment. The specific tools were also chosen to be open source or as accessible as possible to allow them to be made available to the wide range of organizations that make up the crisis response community. The tools support both synchronous activities which take place when team members meet and work together and asynchronous activities when they may work separately and contribute to the knowledge pool that the team is gathering.




Figure 1: Cognitive Work Analysis Phase I – Work Domain Analysis (from Tate et al., 2014)

Community Web Portal

The VCE includes a web-based portal that provides the platform for collaboration and communication, and for creating and sharing resources, as well as more general group-building activity and event awareness (<http://openvce.net> – see Figure 2). After some experimentation and discussion, the open-source Drupal-based software system was adopted as the platform for this site. Drupal is a widely used modular content management system, with an active development community of its own. It provides a user management system and social web functionality such as user profiles, individual blogs and forums. The site was specialised with a range of modules to provide, for instance, twitter-like activity awareness, picture sharing and group management facilities to allow ad hoc teams to be constructed from among the membership as a whole for specific purposes (such as working on a specific response problem). It also includes mechanisms that establish relationships of individuals to the virtual space, allowing users to associate their virtual personae with their real life web profiles. Links are provided to allow users to “teleport” into relevant locations of the virtual world collaboration spaces. This site has been augmented by a wiki (powered by the popular open source MediaWiki software), to provide facilities for co-authoring text documents (a facility felt to be lacking at the time in Drupal).

The screenshot shows the OpenVCE web portal home page. At the top, there is a navigation bar with links for 'About OpenVCE', 'WoSCR Community', 'JOE Community', 'Privacy Policy', and 'Contact us'. The main content area is divided into several sections:

- Welcome:** A banner image with the text: "Welcome to the OpenVCE community portal. All content is community-created, so become a registered user and start contributing!"
- Forthcoming events:** A list of events including "Federal Consortium for Virtual Worlds Conference 2010" and "WoSCR Community - Possible Virtual Iterative Workshop Series - VIWS-4".
- Current discussions:** A section showing recent discussions with user avatars and titles, such as "OpenVCE envisioned site structure" and "The weakest link started by ac".
- What are you doing?:** A section for user activity, featuring a tweet input field and a list of recent tweets from users like Austin Tate and ifuentes.
- 3D space:** A section with a "teleport now" button and access information for chat, wave, HW, and QT.
- Event calendar:** A calendar view for the month of April.
- Right sidebar:** Contains search bars for the site and the wiki, a featured "I-Room" section, "Who's online" status, "Latest wiki updates", and "What's happening" with recent activity.

Figure 2: openvce.net web portal home page

This approach also allows for additional functionality to be made accessible to the community by embedding appropriate tools within site pages. These tools can be generic community tools or introduced for specific tasks such as supporting the team in option generation and pro/con argumentation.

Virtual World Operation Centres – I-Rooms

Warburton (2009) discusses the use of virtual worlds in educational contexts. He provides a table with a rich variety of synchronous and asynchronous communications and presence indication methods, as well as listing some of the issues for usability of virtual worlds like Second Life™ for education and collaboration. These indicate the particular niche which virtual worlds meeting spaces have in providing support to synchronous meeting facilities, as well as showing a large degree of overlap of the issues users have in using such spaces with those found during the OpenVCE project use by the WoSCR community.



Figure 3: A synchronous meeting in a virtual world collaborative space with a range of experimental 3D visualisations for planning and option discussion

Second Life™ and OpenSimulator virtual world environments have been used to realise “I-Rooms” – virtual spaces for intelligent interaction (see figure 3). Using the I-Room concept within virtual worlds gives a collaboration an intuitive grounding in a persistent space in which representations of the participants (their “avatars”) appear and the artifacts and resources surrounding the collaboration can be granted a surrogate reality.

Although for the uninitiated the virtual space can initially be disorienting and video game-like, in our experience users quickly feel comfortable in the space once any technical issues are ironed out (as is the case for other video-conferencing systems, these issues are usually related to audio difficulties or firewalls). Through an avatar a user can see the avatars of other users of the space, and communicate with those in earshot using spatialised voice (communication is also possible using general text chat and instant messaging). This audio-visual positioning in 3D space provides a compelling sense of shared presence with any other users currently in the same space.

In addition to its use as a distributed access meeting space, the I-Room can be used to deliver intelligent systems and tool support for meetings and collaborative activities. In particular, the I-Room is designed to draw on I-X technology (Tate, 2000) which provides intelligent and intelligible (to human participants) task support, process management, collaborative tools and planning aids to participants (Tate et al., 2010a). This technology encourages collaborators to share information about the processes or products they are working on through a common conceptual model called <I-N-C-A> (Tate, 2003). This framework allows access to automated capabilities or agents in a coherent way, providing participants in I-Room meetings with, for instance, access to knowledge-base content and natural language generation technology.

Figure 4 shows an I-Room in a virtual world viewer alongside a browser onto the web portal, typical of how a user's screen might be laid out while using the VCE.




Figure 4: VCE web portal showing virtual collaboration protocol support, alongside an I-Room virtual world operations centre space

Virtual World Collaboration Spaces

The University of Edinburgh created a virtual world collaboration space with a range of facilities that could be provided in platforms such as Second Life™ and OpenSimulator. 3D modelling was provided by Clever Zebra to give professional spaces that could easily be selected or adapted for a range of purposes including:

- A central plaza where arrivals can be directed to appropriate spaces, meetings, events and given news. Users can also pick up (free) items to tailor their avatar or get assistive technologies to improve their access to the virtual collaboration spaces.
- A lecture style seated auditorium with stage space and multiple presentation screens.
- An exposition pavilion to allow for poster displays, demonstration booths and link ups to external web pages and active systems for further exploration. Non-player characters with chat bot capability could be in attendance at the stands when the project team members involved were not present.
- Project and group spaces where artifacts, maps, posters and information boards could persist to allow for small group meetings.
- Home bases and a range of informal meeting spaces to allow for voice and/or text chat discussions without interference with other spaces.
- I-Rooms: virtual spaces for intelligent interaction – an operations centre inspired brainstorming area surrounded by spaces for situation sense making, planning, decision making and communication, command and control purposes. (Tate et al., 2010b)

The collaboration "region" has been packaged as an open educational resource and made available in the Second Life™ marketplace and as a widely available OpenSimulator Archive (OAR) file. Several OpenSimulator-based grids offer the OpenVCE OAR as a starter region when they rent out virtual world space and it has been used in a number of educational environments (e.g., for an oil rig safety training environment, Tait et al., 2017).

MOSES is a virtual world research programme of the United States Army Advanced Training Systems Division (ATSD). The goal is to research and develop advances in virtual world technologies for use in simulation-based training. An OpenVCE region is available for demonstration purposes on MOSES (MOSES, 2017) and the OpenVCE OAR is provided via the MOSES web site for US government users and others to use and adapt to their needs (see figure 5).



Figure 5: OpenVCE Virtual World Collaboration Spaces as distributed in the MOSES OpenSimulator-based Virtual World Platform

As new forms of interaction with 3D and virtual environments evolve, such as using the Oculus Rift and HTC Vive VR headsets, the OpenVCE facilities become potentially even more useful. VR headsets have been used with the OpenVCE collaboration facilities and I-Rooms, especially in OpenSimulator, but also in emerging new platforms such as Linden Lab's Sansar, High Fidelity and Sine Space (Tate, 2015a). See figure 6 for an Oculus Rift “double barrel” virtual reality headset view of the OpenVCE region.




Figure 6: OpenVCE virtual world region as seen via the Oculus Rift virtual reality headset

An “OARConverter” tool (Iseki, 2015) (Tate, 2015b) to take content exported from OpenSimulator as an “OAR” – an OpenSimulator Archive file – via the portable Collada 3D model format into other 3D modelling environments such as Unity3D allows for the OpenVCE assets to be reused in these emerging platforms, extending the utility of the work (Tate, 2016). See, for example, figure 7 which shows the OpenVCE region in the Sine Space (<http://sine.space>) multi-user collaborative virtual world environment.



Figure 7: OpenVCE Virtual World Collaboration Spaces in Sine.Space after conversion from OpenSimulator through Collada and Unity3D to Sine.Space

Simplifications of the Virtual World Collaboration Facility

Later work on the Dismounted Infantry Collaboration Environment (DICE) which aimed to support medics treating injured soldiers and the subsequent reviews of such cases (Tate et al., 2012) led to a simplified and uncluttered version of the virtual world-based I-Room in both Second Life™ and OpenSimulator (see figure 8).

This also was accompanied by a very much simplified entry web page to get users in quickly, and provide direct meeting support in the space. A lot of facilities provided for earlier demonstrations and experimental 3D visualisations in the previous WoSCR I-Rooms were removed to simplify new user engagement and interactions between team members.

The simplified DICE I-Room is included alongside the original I-Room and other collaboration and meeting facilities in the open source distributed versions of the OpenVCE virtual world assets, and in a demonstration environment hosted on the U.S. Army Research Laboratory's own MOSES OpenSimulator-based grid (MOSES, 2017).




Figure 8: Simplified Virtual World Collaboration Space for DICE Project

Experiments and Evaluation

The VCE attempts to facilitate distributed collaboration by integrating asynchronous collaboration through social web technologies and synchronous collaboration through I-Rooms and virtual environments. Two experiments were conducted in 2010 to examine the impact the VCE had on crisis planning and collaboration when compared to traditional means of distributed collaboration among crisis response organizations and individuals. Results and conclusions from the second and more comprehensive of the two experiments are summarised here and reported on in more depth in Tate et al. (2014).

Background. The VCE experiment introduced a biological agent outbreak scenario to two teams of equally staffed crisis expert volunteers distributed across the U.S., U.K., Canada and Italy. The traditional group (control condition) used technology and means that would normally be used for distributed collaboration across these types of organisations (government, industry, non-government, military, and academia) during a crisis, including e-mail for asynchronous collaboration and telephone and teleconferencing for synchronous collaboration. The virtual group (experimental condition) used the full capability of the VCE as described in this paper for synchronous and asynchronous collaboration.

Participants. The virtual group consisted of 10 participants and the traditional group had 7 participants due to 3 individuals who were not available to participate. Each group had what

was considered equal expertise in crisis response and biological outbreaks and had no prior experience working with each other. The groups each had at least one international member and had representatives from government, academic, non-profit, and industry organisations. Each group was given the same scenario and asked to generate a crisis response plan over four days.

Communication Patterns. The distribution and quantity of e-mails for each group was analysed as an indicator of their patterns of asynchronous communications. The basic pattern of most e-mails being addressed to the entire group was similar across both groups. However, the virtual group sent 38% fewer e-mails through the 4-day planning effort compared to the number of e-mails sent by the traditional group. A large difference in the patterns of communication was that the virtual group made heavy use of the online portal pages, forum, and wiki capabilities that produced a total of 2098 total page views or visited within the site. These results indicate that the asynchronous communications were different and that the virtual group did in fact use the portal capabilities for communication though they had the choice of using only e-mail for their asynchronous communication needs.

Uncertainty. Among one of the measurements taken each experiment day was a measure of uncertainty for each participant. Uncertainty was evaluated along two dimensions, namely goal and procedural uncertainty. Goal uncertainty is defined as the level of ambiguity a person has about the goals or objectives in their current situation or task. Procedural uncertainty, on the other hand, is how much ambiguity is associated with the steps or procedures necessary to accomplish the defined goals. Two seven-point Likert scale items measured each uncertainty dimension, which were averaged together. Choo (2005) has defined these uncertainty dimensions in terms of their interactions with each other. The amount of goal and procedural uncertainty possessed by an individual and group will dictate the mode (see Figure 9) of interactions and ultimately the success of the group.




Figure 9: Goal and procedural uncertainty dimensions and the various modes of interaction they can create based on the levels of uncertainty for each dimension

Placing the results for goal and procedural uncertainty along the uncertainty dimensions presents a clear picture of how much uncertainty was involved for each group (see figure 10). The traditional group finds themselves interacting in the “anarchy mode” where there is ambiguity with both goals and procedures. Group and individual feedback after the experiment confirms this finding. There was considerable effort needed by this group to establish a common ground and understanding within the group before they could engage in any planning efforts. This is also indicative of collaboration efforts among many different organizations, involving people with different backgrounds and expertise, particularly when they have not worked together before. The virtual group using the VCE and virtual collaboration protocol fared much better and found themselves working within the “relational mode” where goals and procedures are clear and understood. The overall difference between the two groups was statistically examined using repeated measures analysis of variance (ANOVA) and there was a significant difference between the two groups as suggested in Figure 10 ($F(1, 15) = 10.31, p < .01$). The virtual group had less goal and procedural uncertainty as they collaborated with their colleagues, which can result in increased efficiency and performance.




Figure 10: Goal and procedural uncertainty results from the virtual and traditional groups showing differences in uncertainty and the mode each group was interacting within

Time on task. The participants were asked to estimate the total time spent on planning efforts outside of their synchronous meetings. The self-reported average durations for the virtual group was 5.2 hours compared to 3 hours for the traditional group. An independent

samples t-test showed that this difference was significant ($t(8) = 2.88, p < .05$). The virtual group has shown that they are communicating asynchronously differently than the traditional group and these results indicate that they spent more overall planning time as they were interacting with their teammates.

Planning output. The final assessment area analysed the quality of their planning efforts. The participants were instructed to generate a planning document based on the crisis they were given. Their planning documents were analysed with text mining software called Leximancer (www.leximancer.com). This automated content analysis software was used to identify the semantic concepts addressed by each of the groups (Smith & Humphreys, 2006). Part of the output is a visual concept map that indicates the concepts in the document s produced and their relationships (figure 11).




Figure 11: Concept map that indicates automated content analysis of the semantic concepts addressed by each of the groups in the planning document they produced

In figure 11, both team plans can be seen (“FILE_virtual team plan” and “FILE_traditional team plan”) along with a “bio outbreak referent” which was based on a documented procedure for a biological outbreak response. The virtual team plan shows more depth due to its four originating concept links (flu, reindeer, risk, and response – each one level deep from the document root) versus the one concept link found with the traditional plan (medical). The virtual plan also displays greater breadth due to the larger number of concepts included in its plan versus the traditional plan. Finally, when each plan is compared to the bio outbreak referent plan, the virtual team shares a common concept group (vaccine) with the referent plan while none are shared with the traditional plan. These results suggest that the virtual group’s plan addressed more concepts, developed those concepts in more detail, and addressed concepts more similar to that of documented procedures compared to the plan the traditional group generated.

Results summary. The results when combined provide a possible explanation to the higher quality plan produced by the virtual group. The planning process begins and is facilitated through communication. The patterns of communications between the groups were shown to be quite different with the virtual group augmenting their e-mail communication extensively with the Web 2.0 capabilities found in the provided portal. These interactions by the virtual group resulted in less procedural and goal uncertainty, which could be the reason why the virtual group members chose to spend more time on their planning tasks compared to the traditional group members. If the end-goal is clear and the way to achieve that goal is clearly defined, it is much easier for individuals to spend time toward that task. Given that the virtual group spent more time on a task that they defined more clearly by means of their enhanced communication process, it is not surprising that they produced a higher quality plan in the end.

Conclusion

The Virtual Collaboration Environment (VCE) is based on a combination of Web 2.0 community knowledge sharing and collaboration tools which can be used asynchronously and a shared online virtual meeting space which can allow for synchronous meetings of the distributed team. A virtual collaboration protocol supports the activities of the group to improve the value of the environment to the participants.

A Cognitive Work Analysis was used to refine the requirements for a target community of experts engaged in support to large scale emergency crises and specific tools were selected to meet their needs. Experiments were performed in a number of emergency response simulations (based on a biological disease outbreak in a large city with input from an international team of experts) to show that there were improvements in Goal (Objective) definition and procedural (Decision and Action) outcomes when the virtual collaboration environment was used when compared to traditional telephone and teleconferencing support. Higher quality plans were produced by the team supported by the virtual collaboration environment.

The virtual world collaboration space, meeting support tools and associated content have been provided to allow others to use and customise the facilities for their own needs. The resources have been released under flexible open source as the Open Virtual Collaboration Environment (OpenVCE).

Acknowledgements

The OpenVCE project was funded by U.S. Joint Forces Command, the Army Research Lab and Alion Science and Technology Corporation. Specific applications mentioned in this article have been sponsored by a number of organizations. The University of Edinburgh, U.S. Army Research laboratory and research sponsors are authorized to reproduce and distribute reprints and online copies for their purposes notwithstanding any copyright annotation hereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of other parties. Drupal is a registered trademark of Dries Buytaert and Second Life™ is a registered trademark of Linden Research, Inc. The authors acknowledge the contribution to the work presented herein of their colleagues from University of Edinburgh, the University of Virginia, Carnegie Mellon University, Perigeon Technologies LLC and Clever Zebra, and that of the many volunteers and professional emergency responders who generously donated their valuable time and expertise to take part in the experiments.

References

- Choo, C. H. (2005). *The Knowing Organization: How Organizations Use Information to Construct Meaning, Create Knowledge, and Make Decisions*. Oxford University Press, New York, NY.
- Cross, R. and Parker, A. (2004) *The Hidden Power of Social Networks: Understanding How Work Really Gets Done in Organizations*, Harvard Business School Press.
- Hansberger, J.T., Tate, A., Moon, B. and Cross, R., *Cognitively Engineering a Virtual Collaboration Environment for Crisis Response*, Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Working. (CSCW 2010), Savannah, Georgia, U.S.A, 6-10 February 2010. Abstract and Poster.
- Iseki, F. (2015) OARConv. <http://www.nsl.tuis.ac.jp/xoops/modules/xpwiki/?OARConvWin>
- Moon, B., Hansberger, J.T. and Tate, A. (2011) *Novakian Concept Mapping in Virtual Collaboration Environments*. In: B. Moon, R.R. Hoffman, J. Novak and A. Canas (Eds.) *Applied Concept Mapping – Theory, Techniques, and Case Studies in the Business Applications of Novakian Concept Mapping*, CRC Press.
- MOSES (2017) *MOSES Military Metaverse – Downloads – Content – OpenVCE OAR*. <https://militarymetaverse.org/content>

Smith, A.E. and Humphreys M.S. (2006). Evaluation of Unsupervised Semantic Mapping of Natural Language with Leximancer Concept Mapping. *Behavior Research Methods*, 38 (2), 262-279.

Tait, J., Hetherington, C. and Tate, A. (2017) Enhancing Student Employability with Simulation: The Virtual Oil Rig and DART, Poster Presentation, 3rd International Enhancement in Higher Education Conference: Inspiring Excellence - Transforming the Student Experience, 6th-8th June 2017, Radisson Blu Hotel, Glasgow, UK. The Quality Assurance Agency for Higher Education, UK.

Tate, A. (2000) Intelligible AI Planning. In *Proceedings of ES2000, The Twentieth British Computer Society Special Group on Expert Systems International Conference on Knowledge Based Systems and Applied Artificial Intelligence*, pp. 3-16, Springer.

Tate, A. (2003) <I-N-C-A>: a Shared Model for Mixed-Initiative Synthesis Tasks. In *Proceedings of the Workshop on Mixed-Initiative Intelligent Systems (MIIS) at the International Joint Conference on Artificial Intelligence (IJCAI-03)*. Acapulco, Mexico, August 2003.

Tate, A. (2006) The Helpful Environment: Geographically Dispersed Intelligent Agents That Collaborate, Special Issue on "The Future of AI", *IEEE Intelligent Systems*, May-June 2006, Vol. 27, No. 3, pp 57-61. IEEE Computer Society.

Tate, A. (2015a) OpenVCE in Unity and VR, Blog Post, 5th October 2015.
<https://blog.inf.ed.ac.uk/atate/2015/10/05/openvce-in-unity-and-vr/>

Tate, A. (2015b) OpenSim OAR Convert to Unity Scene with Windows Interface, Blog Post, 24th October 2015. <https://blog.inf.ed.ac.uk/atate/2015/10/24/opensim-oar-convert-to-unity-scene-with-windows-interface/>

Tate, A. (2016) Sine.space – OpenVCE Region, Blog Post, 15th November 2016.
<https://blog.inf.ed.ac.uk/atate/2016/11/15/sine-space-openvce-region/>

Tate, A., Dalton, J., Stader, J., Wickler, G. and Hansberger, J.T. (2007) Collaborative Operations for Personnel Recovery Final Report on DARPA/AFRL Contract No. F-30602-033-2-0014, 31-Aug-2007 (revised 30-Jan-2008).

Tate, A., Chen-Burger, Y-H., Dalton, J., Potter, S., Richardson, D., Stader, J., Wickler, G., Bankier, I., Walton, C. and Williams, P.G. (2010a) I-Room: A Virtual Space for Intelligent Interaction, *IEEE Intelligent Systems*, Vol. 25, No. 4, pp 62-71, July-August 2010, IEEE Computer Society.

Tate, A., Chen Burger, Y-H., Dalton, J., Potter, S., Wickler, G., Carley, K.M., Kunkel, F., Cross, R., Hansberger, J.T. and Moon, B. (2010b), Open Virtual Collaboration Environment for the Whole of Society Crisis Response Community, *Proceedings of the Sixth International Conference on Knowledge Systems for Coalition Operations (KSCO-2010)* (Lawton, J. (ed.)), Vancouver, B.C., Canada, September 21-23, 2010.

Tate, A., Hansberger, J.T., Potter, S. and Wickler, G. (2012), Support for Distributed Collaboration in the Dismounted Incident Collaboration Environment (DICE), Proceedings of the Seventh International Conference on Knowledge Systems for Coalition Operations (KSCO-2012) (Lawton, J., Patel, J., Suri, N. and Tate, A. (eds.)), Pensacola, FL, U.S.A, February 15-16, 2012.

Tate, A., Hansberger, J.T., Potter, S. and Wickler, G. (2014) Virtual Collaboration Spaces: Bringing Presence to Distributed Collaboration, Journal of Virtual Worlds Research, Assembled Issue 2014, Volume 7, Number 2, May 2014.

Vicente, K. J. (1999). Cognitive Work Analysis. Lawrence Erlbaum Associates, Mahwah, NJ.

Wickler, G., Tate, A. and Hansberger, J.T. (2007) Supporting Collaborative Operations within a Coalition Personnel Recovery Center, Proceedings of the Fourth International Conference on Knowledge Systems for Coalition Operations (KSCO-2007) (Lawton, J., Patel, J. and Tate, A. eds.) pp. 14-19, Waltham, MA, U.S.A, 1-2 May 2007. In Proceedings of the International Conference on Integration of Knowledge Intensive Multi-Agent Systems Modelling, Evolution and Engineering (KIMAS '07).

Wickler, G., Tate, A. and Hansberger, J. (2013), Using Shared Procedural Knowledge for Virtual Collaboration Support in Emergency Response, in Special Issue on Emergency Response, IEEE Intelligent Systems, Vol. 28, No. 4, July/August 2013, IEEE Computer Society.

Knowledge Systems for Coalition Operations KSCO 2017

November 6-8, 2017
Los Angeles, California, USA

A Semantic-Based Information Integration Framework of Agile Command and Control

Paper ID Number:

72

Authors:

Wen Zhou(National University of Defense Technology)
Weidong Bao(National University of Defense Technology)
Xiaomin Zhu(National University of Defense Technology)
Chao Chen(National University of Defense Technology)
Zhu Wu(Naval Command College)
Changfei Wu(Air Force Early Warning Academy)

Point of Contact:

Wen Zhou
zhouwen@nudt.edu.cn

Relevant Topic Areas:

Intelligent Command and Control (C2)
Coalition and Team Information Sharing
Coordination and Collaboration

A Semantic-Based Information Integration Framework of Agile Command and Control

Wen Zhou^a, Weidong Bao^a, Xiaomin Zhu^a, Chao Chen^a, Zhu Wu^b, Changfei Wu^c

^a*College of Information Systems and Management, National University of Defense Technology, Changsha, China;* ^b*Naval Command College, Nanjing, China;* ^c*Air Force Early Warning Academy, Wuhan, China*

Email: zhouwen@nudt.edu.cn

Abstract: Human society is in the period of accelerated development of the information revolution. It not only affects people's life extensively, but also constantly updates our understanding and understanding of the mode of operation in the information age. The ability to carry out effective information integration and provide intelligent information services, has become a joint combat process to obtain the command and control of the agile advantage of the support base. Based on this, this paper proposes a semantic-based information integration framework of agile command and control, which includes two parts: basic module and semantic analysis engine. The basic idea is to change the semantic relation from the pursuit of comprehensive information reasoning to the attention data, the idea of information retrieval into knowledge retrieval, relying on the information integration framework to tap the command information system implied in a variety of rich semantic relations, in order to carry out a higher level of information integration and provide more intelligent knowledge retrieval, and thus agile Command and control.

Key words: agile command and control; information integration; semantic analysis; graph database

1 Introduction

On the battlefield, to carry out agile distributed command and control(C2) for a force composed of a joint force of different countries or different arms, the first problem to be solved is how to realize the interoperability between the command information systems used by different command centers, and the core is the effective integration of information^[1]. Besides, agile C2 also requires the command information system to achieve more intelligent knowledge retrieval, rather than the information retrieval has been.

With the development of Internet technology and the wide application of service-oriented technology, many information systems have gradually begun to integrate information according to service needs, and semantic information integration has also made great development. However, in order to achieve agile C2, we also need a higher level of information integration, to achieve more intelligent knowledge retrieval, and these are still some fundamental difficulties.

The first is the limitations of relational databases. Developers have been trying to use relational databases to deal with associated or semi-structured data sets, such as Vysniauskas^[2] proposed a hybrid approach for correlating OWL 2 ontologies and relational databases, Chhaya^[3] proposed to use D2RQ and Ontop to publish relational

databases as associated data, which is a framework for answering SPARQL queries in relational databases^[4]. Similar to the Ontop framework, there is GRAPHITE, which is a framework for implementing extensible graph traversal in relational database management systems^[5], and Jindal^[6] also proposed the use of vertical relational databases for graph analysis. However, as the outliers increase, the macroscopic structure of the data set becomes more complex and irregular, and the relational model will result in a large number of table joins, sparse rows, and nonempty checking logic. Increased connectivity in the world of relationships will translate into an increase in connection operations that will hamper performance and make it difficult for existing databases to respond to changing business requirements. Whether trying to model or correlate associations in a relational database, in addition to increasing the complexity of queries and calculations, it is necessary to deal with the pattern of the double-edged sword, and many times the pattern proved too rigid and fragile. In response to this problem, Cerans^[7] proposed semantic refactoring of relational databases, resulting in a modeling framework for semantic databases such as SpiderMass^[8], which combined with the social needs of social networks Extensive development, including semantic networks^[9] and graph databases^[10], such as Graphx^[11], Mizan^[12] and TAO^[13] are very good graph database management systems.

Secondly, the problem of semantic databases defined by ontology modeling language. In order to improve the interoperability between different systems, such as in the standard data exchange format, ontology and consensus-based information model have made great progress. In terms of theory, not only the study of command and control ontology, but also Singapogu discussed the role of ontology in C2SIM^[14], and Hansen proposed an information integration scheme based on ontology matching in time constraints^[1]. In terms of applications, it includes the use of probabilistic ontology modeling methods to design terrorist decision-making support systems^[15], the development of ontology on the hypothesis management of sea development^[16], and the mission and means to achieve military assets and mission objectives^[17]. However, while ontology is widely accepted, different information systems adopt a variety of different ontology modeling languages, such as Web ontology language, where the interaction between them and the different grammar rules will cause the developer to pay a high cost^[18]. And, in order to cope with the adaptability of data requirements, which need to build the semantic library through the ontology modeling language must develop a variety of applications and interfaces^[19], where the definition of mapping rules and application re-design and development is a very tedious work^[20].

Finally, in order to achieve the perception of the user's intention and enhance the intelligence of the information system, one of its core is the understanding of semantic information system, and there are still a lot of technical problems need to be resolved, such as natural language processing, semantic and artificial Intelligent and so on. In terms of natural language processing, including word sense disambiguation^[21, 22], identification of entities and relationships^[23], extraction of syntactic and semantic features^[24-26]. In semantics, including ontology-based semantic search engine development, common sense and semantic reasoning^[27-29] and so on. And the key is to change the concept of information retrieval, that is, by the information retrieval and

keyword matching into knowledge retrieval, and the use of semantic and artificial intelligence to achieve self-learning^[30], but for now we need to do more in-depth study.

In order to improve the interoperability of information systems and achieve a higher degree of agile C2, we propose a semantic-based information integration framework of agile C2 for the higher level of information integration of command information systems. We construct the middleware by using the Extract Transform Load(ETL) method based on subgraph to realize the transformation from the relational database to the graph database. The core is to realize the deeper and more efficient processing of the associated data by digging the data connection. In order to provide the domain knowledge support for the search query of the database, we also develop the semantic analysis engine to realize the knowledge retrieval based on semantics and further improve the intelligence of the retrieval, and then improve the retrieval of the database, so as to provide users with more intelligent, personalized and professional information services.

The rest of this article is arranged as follows. In Section 2, we introduce a semantic-based information integration framework of agile C2. Section 3 describes the basic modules in the information integration framework, including the relational database to map database conversion, domain model building and syntax analysis. Section 4 focuses on the core of the information integration framework -- the composition of the semantic analysis engine, and an example is described. Section 5 is the conclusion.

2. Semantic-based information integration framework of agile C2

In the process of joint operations, in order to achieve a higher degree of agile C2, we not only need to achieve interoperability between different information systems, and we also need information systems to understand the user's intent to achieve semantic-based knowledge retrieval, and further enhance the retrieval of intelligent to provide users with more intelligent, personalized and professional information services. At present, to achieve this goal also has many difficulties and challenges, especially information integration is the foundation. To solve this problem, we propose a semantic-based information integration framework of agile C2, as shown in Fig. 1.




Fig.1 Overall system architecture

In terms of composition, we include two basic parts of the basic module and the semantic analysis engine in our proposed framework. Among them, the basic module includes middleware, a graph database, a domain model and a parsing module for converting the relational database into a graph database, and the semantic analysis engine includes a universal package module, an extension semantic module, a matching definition module and a search module constitute.

In terms of process, our framework mainly includes four processes. First, we develop the relational database into the graph database through the middleware. Secondly, we develop the domain model and establish the mapping rules of the domain model and the graph database. Then we introduce the syntax parsing module of natural language processing, and implement the structured processing of the query. The Finally, we develop a semantic analysis engine to achieve semantic-based knowledge retrieval.

3. Basic module

3.1 Middleware and graph database

With the advent of large data age, NoSQL database has seen rapid development. Graph Database is one of the NoSQL family and one of the most developed database technologies since 2013. Common database systems include AllegroGraph, DEX, HyperGraphDB, InfiniteGraph, Neo4j, and so on^[10]. Some relational database management systems also began to support graph data, such as Oracle to increase the space of large data and data support. It can be said that the database database technology is one of the most popular research topics.

In the aspect of multi-level relationship, shortest path, PageRank and so on, the graph database can adopt the matching algorithm of graphs, and the iterative level is less in the node relation query process. The query efficiency is obviously better than the relational database, especially when the data is large, the advantage is more obvious.

Although the semantics can be described to a certain extent by the graph, and the description and understanding of semantics occupy an important position in the C2 process of joint operations. However, the existing data is stored in the form of relational data, In the form of organizational data, between the table through the primary key - foreign key for the association, the way is simple. In contrast to this, the graph data is stored in the form of graphs, nodes and edges are the basic representation elements of the graph, and the data representation is more complex^[31]. Therefore, migrating from a relational database to a graph database is a practical solution. How to achieve the relationship between the data to the map data conversion is to achieve the existing application to the data application transformation key. The problem to be solved in this paper is how to improve the quality of the converted graph data in the process of ETL relational data to graph data, and to convert the relational data into graph data efficiently and efficiently.

ETL is actually a process of describing data extraction, conversion, and loading of data from source data. Common tools include Informatica, Datastage, OWB, MS DTS, Beelod, Kettle and so on^[31]. ETL quality is mainly measured from the aspects of

correctness, completeness and consistency. Because traditional ETL tools such as Informatica and Kettle are mainly used for the research and practical application of relational data to relational data extraction, transformation and conversion from relational data to relational data using ETL. However, in the use of ETL to achieve from the relational data into graph data, due to the relatively late development of the database, and the standard is not uniform, so the research and application is relatively small. Now the main problems include the quality of the converted graph data is not high, when faced with complex relational database conversion efficiency is not high, the conversion results are not conducive to distributed storage.

Therefore, by transplanting some successful algorithms in traditional ETL, such as drawing some successful methods of ETL under the framework of Spark GraphX, we adopt a method of efficient batch data extraction, parallel conversion and batch loading. GraphX is the development of Apache distributed graph calculation framework, which is mainly to solve the problem of distributed computing. It provides ETL, which provides data extraction, cleaning, conversion, loading tools and exploratory analysis, and through a single system for iterative calculations. In addition, it can view the same data, graphics and collections, through RDDs efficient transformation, connecting graphics, and can be customized through the Pregel API iterative algorithm. It is comparable to the fastest professional-grade graphics processing system while preserving flexibility, fault tolerance and ease of use.

In this paper, we use the ETL method based on subgraph proposed in [31] to construct the middleware, including six steps.

(1) Load relational database schema: read the relational database schema definition information, its table definition, the relationship between the definition were loaded into the corresponding list.

(2) Establish the relational database model corresponding to the model diagram: that is based on the relational database table definition, relationship definition information to build the relationship model.

(3) Split the relationship between the database model for a number of sub-mode: traversal mode diagram, get the corresponding set of points C , according to the relationship between the model will be split into several sub-mode.

(4) Loop processing sub-mode: parallel pairs of sub-mode ETL, according to the sub-mode of the main table definition information, bulk loading of the main table records, the relationship between these records to map data conversion. And then according to the sub-mode of the relationship between the table, followed by loading the other table data, and the relationship between the data to the map data conversion. Finally, we get the subpattern pattern and subgraph data corresponding to the subpattern.

(5) Optimize the conversion results: the use of indexing, merging nodes, mergers and relationships, etc., to reduce the complexity of graph data and improve the efficiency of data query.

(6) Store the conversion results, the conversion results will be stored in bulk to the database.

In addition, distributed storage can be used as a choice for large scale data storage management. In this paper, the distributed graph data analysis method is based on the

distributed storage of graph data. Through the design and management of control procedures, optimizing and integrating the database storage technology and graph data distributed computing framework, we finally unified the distributed data storage and distributed computing.

3.2 Domain model

In order to construct the domain model, we refer to the idea of the ontology architecture, as shown in Fig. 2, in the reference [18, 32]. We use not only domain experts to build domain models, but more importantly, we use the common Joint C3 Information Exchange Data Model(JC3IEDM)^[33], and combined with different tasks to guide the domain model to create different domain ontology. Moreover, in order to make the domain model of each component can be compared with each other, we allow each component has its own semantic description on the basis of each component is required from the shared ontology. This approach allows us to overcome the comparative problems in multiple ontology methods and also alleviate the shortcomings of a single ontology approach that is susceptible to changes in source information.

In addition, our global shared ontology comes from the open knowledge base, such as DBpedia, Knowledge Graph, Wikidata, Freebase and so on. They are an important complementary to JC3IEDM and other data models, whose main role is to improve the command control information integration process of data interoperability. Take DBpedia as an example. DBpedia's concepts and relationships to facilitate data interoperability in command information integration. DBpedia is Wikipedia, which is a knowledge base of human knowledge resources, maintained by a large number of people. The project balances this huge knowledge resource by extracting structured data from Wikipedia's entries in a more efficient way. Based on Wikipedia datasets, DBpedia allows users to query complex issues and link other data sets to other Wikipedia datasets. At present, DBpedia describes more than 340 million events, of which 1.5 million have been classified as ontology.




Fig.2 Domain model architecture

It should be noted that DBpedia storage format is RDF triple, and RDF is proposed by the W3C to describe the network resources standards. It uses a simple description of

the way, that is, the subject (subject), predicate (predicate), object (Object) composed of triples to represent the resources. The RDF data general representation is a (S, P, O) triplet, and a set of RDF data can form a RDF directed graph. The RDF graph can be represented by a tagged node and a tagged edge, where each triple corresponds to a "node-edge-node" subgraph on the graph and it states the relationship between the subject and the object of the object as expressed by the predicate. The node of a RDF graph is the subject and object of all the triples it contains, and the direction of the edge always points to the object. RDF graphs can usually be viewed as a directed marker.

In order to facilitate the construction of the domain model, and to achieve docking with the map database, we use the map database Neo4j on the collection of DBpedia RDF triples stored. Neo4j is called a property graph, including vertices, relationships, and properties^[34]. The storage of data is mainly divided into three types of data storage, such as nodes, relationships, nodes or relational properties. Whether it is a vertex or relationship, can have any number of properties, the property is similar to a hashMap storage. Neo4j focuses on the performance of a large number of connection queries when the performance of the problem. In addition, Neo4j also provides a very fast graphical algorithm, recommended system and OLAP style analysis, it can be said, Neo4j is a high performance, high reliability, scalable, fully compatible ACID map database^[34]. Since each Node, Relationship, and Property in Neo4j are stored independently and follow the natural order, if you want to find the corresponding node in the diagram, you must rely on the index. Neo4j mainly provides Lucene-based Full-text index mechanism to achieve the search for nodes and relationships. The RDF data storage model diagram based on the graphics database Neo4j is shown in Fig. 3.




Fig.3 RDF data storage model based on graphic database Neo4j^[34]

3.3 Syntax parsing module

The syntax parsing module is the basis of semantic retrieval. After receiving the user to retrieve the task, the first need for the user's search task for syntactic analysis, which belongs to the natural language processing field of information extraction problem. The joint extraction of entities and relations as a key task of information extraction, the implementation method can be simply divided into two categories: one is the series extraction method. The other is a joint extraction method. The concatenation method divides the problem into two sub-tasks in series, that is, the entity

extraction model is used to extract the entity, and then the relational extraction model is used to get the relationship between the entity pairs. The advantage is that it is easy to optimize the entity recognition task and the relationship extraction task, but the disadvantage is that they are aimed at obtaining the intermediate product (entity or relationship type) of the triplet, and the result of the entity recognition will further affect the result of the relationship extraction, resulting in error accumulation^[35]. Unlike the concatenation method, the joint extraction method uses a model to extract entities and their relationships simultaneously, enabling better integration of information between entities and their relationships. However, there are many problems with the existing joint extraction methods, such as: most of the joint extraction model requires manual participation in the construction of features; based on the end to end of the joint extraction model, because in the process of modeling the extraction of entities and their relationship between the information Redundancy and other issues^[35].

In order to realize the automatic labeling of the sequence of text words, we use the end-to-end model proposed in [35]. The model structure is shown in Fig. 4. In this model structure, it contains a bi-directional long memory (Bi-LSTM) layer for encoding the input statement and the LSTM-based decoding layer with offset loss. Deviations can increase the relevance of physical labels. The word embedding layer transforms the one-hot representation vector of each word into a low-dimensional dense word embedding vector (dimension 300), and the Bi-LSTM encoding layer (the number of layers is 300) is used to obtain the coding information of the word, the LSTM decoding layer (the number of layers is 600) is used to generate the tag sequence. We add offset loss to enhance the relevance of the entity tag.




Fig.4 An End-to-End model to produce tags sequence^[35]

4. Semantic analysis engine

4.1 Semantic Analysis Engine Architecture

The ultimate goal of developing a semantic analysis engine is to achieve semantic-based knowledge retrieval. Semantic analysis engine is the core component of our proposed semantic-based information integration framework of agile C2, which includes five groups, including crowdsourcing module, extension semantic module,

external reference library, matching definition module and search module, as shown in Fig. 5.




Fig.5 Semantic analysis engine architecture

The semantic analysis engine's workflow is as follows:

The first is to expand the semantic module. As the syntax analysis module has extracted the entity and the relationship from the user's query task, in the semantic analysis engine, we need to further extract the extracted entities and relationships as the extended labeled property graph. Its main function is to enter the subgraph as a matching definition. In the process of modeling the extended labeled property graph, it requires the user to enter or define a template, which is mainly used to define the query question and focus of the user. At the same time, we can also enter or identify property information related to a specific field in the template, and this information is recorded in the log.

The second is an external reference library. It consists of open knowledge base, domain model and JC3IEDM, in which open knowledge base provides a wide range of factual knowledge, and domain model and JC3IEDM provide C2 domain knowledge. Because the extended labeled property graph needs to be guided with task-related knowledge in the process of modeling, and it also needs to be provide constraints in the search process, the external reference library can provide extended knowledge and constraints as effective guidance.

The third is the crowdsourcing module. It is mainly based on the log for the next

step in the definition of matching to provide more information available, which is to further realize the semantic retrieval of an important basis for intelligence. In the crowdsourcing module, we will record and organize the system in the user's query records and feedback, on the other hand we will record the user's organizational structure and other related information and label information. This information can provide assistance for more accurate matching of user search intentions. In addition, the template information of the user query and the feedback to the final search result are automatically recorded in the log.

The fourth is the matching definition module. Its main function is to define the subgraph and ranges that need to be matched for the next search, with the purpose of providing a context for the search. It aggregates the information in the crowdsourcing module and the extended semantic modules and converts them into subgraph with matching breadth and depth constraints that need to be matched.

The last part is the search module. The main function of this module is to connect the semantic analysis engine and the graph database, its goal is to convert the subgraph with search breadth and depth constraints into the query language of the graph database and search the database in the graph database.

In the semantic analysis engine, how to model the expansion of the labeled property graph is the foundation. Here we focus on the process of the proposed labeled property graph.

The extended labeled property graph is based on the labeled property graph, and it is the concept of the graph model in Neo4j. The labeled property graph is made up of nodes, links, properties, and labels. Among them, the node contains properties that can exist in any form of key-value pairs. In Neo4j, the key is usually a string, the value can be Java string and primitive data, or an array of these data types. The nodes can be tagged with one or more tags, which label the nodes together and represent their roles in those data^[34]. Links connect nodes, thus forming the graph. Each link has a direction, a name, a start node, and an end node. The direction and name of the link make the structure of the node rich in semantics. And the link can also have properties, by adding properties on the link, you can provide metadata to graph algorithms, also can add additional semantics (including features and weights). In addition, it can also be used for runtime constraint query. Because the labeled property graph can provide these advantages of semantics, our modeling process with it as the core foundation. However, since the properties of the nodes and links in the labeled property graph exist only as semantic metadata, the semantic metadata itself has a very rich intrinsic relationship, especially when combined with user intent and contextual information. Therefore, in order to make full use of the content of these semantic links, we expand the labeled property graph, relax the constraints on the properties of nodes and links, and establish links between nodes and associated properties based on user intent and contextual information.

An example of modeling the expanded labeled property graph is shown in Fig. 6. The user's input template provides user intent information, the external reference library provides the basis for property association and label selection, and the crowdsourcing module provides contextual information for the extended labeled property graph

modeling. It should be noted that the expanded labeled property graph contains a mechanism of learning mechanism, that is, through structure learning, property learning, path learning, rule learning and multi-modal learning to establish an embedding learning mechanism in the process of modeling. And then we can use this self-learning process to continuously strengthen the function of the expanded labeled property graph.




Fig.6 An example for the modeling of extended labeled property graph

In Fig. 6, we firstly need to model the user intent information according to the modeling process of the labeled property graph in Neo4j. There are three nodes A, B, and C, and each node has three properties information. Then we use the external reference library and crowdsourcing module with related knowledge and contextual information to provide more links between the properties of the nodes and links in the labeled property graph. For example, we can establish an association between the B-2 property information of the Node B and the A-3 property information of the A node.

4.2 Example

As an example, we assume that there is a user's input template which actually represents a semantic search scenario for resource in a joint combat C2 process: At T + K, the troop A arrives at the C zone via the B zone and, in conjunction with the troop E, strikes against the combatant D, and the troop E locates in the F zone.

How do you semantically model the C2 intent reflected in the semantic search scenario for resources? According to the information integration framework proposed in this paper, the core is to design the C2 intention of the user input template using the extended labeled property graph, and then provide a richer implicit semantic information service. For example, it can answer the question in the user's input template that what impact factors affect the troop A with the troop E to implement combat operations? That is, in the troop A and E operations, what factors that affect collaboration they need to consider (such as regional location, level of equipment, combat effectiveness, environmental factors, etc.)? And what is the relationship between these factors?

We first express the C2 intent as:

- Troop A passes through zone B.
- Troop A arrive at the C zone.
- Troop A is with Troop E.
- Troop A strike at the combatant D.
- Troop E strike at the combatant D.
- Troop E locates in the F zone.

Fig. 7 shows the basic modeling results for the above C2 intent.




Fig.7 An example for semantic search (the basic modeling result)

On the basis of Fig. 7, we need to use the external reference library and the crowdsourcing module to exploit richer C2 semantic information. The external reference library provides the basis for property association and label selection, the crowdsourcing module provides contextual information for the extended labeled property graph modeling, and the user's input template also provides questions and focus about the C2 intent and other relevant information. We use the expanded labeled property graph to model. That is, we link the properties and select labels according to the C2 domain knowledge in the domain model and JC3IEDM in the external reference library. And use the contextual information in the crowdsourcing module and the open knowledge in the external reference library to automatically add more property information and potential links between different properties for the entities.

It should be noted that our model relies on the user input template, domain model and JC3IEDM in the external reference library to determine the focus of the model and the breadth and depth of the relationship extension, thus forming a final available subgraph for searching and performing a search in the graph database. Fig. 8 is the final result of modeling using the extended labeled property graph proposed in this paper.

5 Conclusion

The joint operation of the information age requires C2 to be agile, and effective information integration is the basis for supporting the C2 agility^[32], and it further requires the command information system to provide intelligent information services such as knowledge retrieval. Although so far there have been many different solutions, including the use of semantic web ideas in the process of information integration to achieve reasoning function. But we found in practice that there are still many difficulties in implementing the reasoning in the process of information integration. Therefore, we try to set our goal to carry out a higher level of information integration, to provide more

intelligent knowledge retrieval, and then to achieve agile command.




Fig.8 An example for semantic search (the final modeling result)

To this end, this paper presents a semantic-based information integration framework of agile C2. The core includes two aspects. First, we convert the relational database into a graph database to facilitate semantic-based knowledge retrieval. The second is that we construct the crowdsourcing module and the extended the semantic module in the semantic analysis engine. Moreover, we introduce an external reference library and use the open knowledge and C2 domain knowledge in the external reference library, together with the contextual information provided by the crowdsourcing module, to explore the richer semantic links from the C2 intent, and to provide users with more intelligent, personalized and professional information service.

It should be noted that, according to the classification method proposed in the reference [36], our goal is to establish a framework that can integrate data model, expression rule, knowledge reasoning and semantic retrieval from the perspective of vocabulary, grammar and ontology. This framework has adaptability, extensibility and versatility. We hope that this framework will facilitate communication, address the ambiguity of needs and understanding, and provide help for the proper realization of semantic-based agile C2 information integration. However, there is no in-depth consideration of agreements and services in our proposed framework, and implementation and testing tools are not yet fully implemented, and these are an important part of our future work.

In addition, in our framework, the external reference library plays an important role,

especially JC3IEDM, which provides not only the support of the establishment of the C2 domain model, but also the basis of liking properties and selecting labels for the extended labeled property graph. However, JC3IEDM is stored on a traditional relational database which doesn't support semantic retrieval well, and it doesn't contain rich semantics. Moreover, it can't represent all the terms in the field of operations and cover the whole relationship among them^[37]. Therefore, we support good semantic retrieval through graph database, and further enrich the semantic information in our model through the contextual information provided by open knowledge base and crowdsourcing module. One of our future work is to improve the semantic analysis engine in the C2 information integration framework by establishing a more effective context-aware mechanism, improving the self-learning mechanism of the extended labeled property graph, and regulating the mechanism of user input template.

Acknowledgements

This work is supported by National Natural Science Foundation of China (NSFC No.71702186), the Scientific Research Project of National University of Defense Technology under grant ZK16-03-57. We thank all anonymous reviewers for their constructive comments.

References

- [1] B. J. Hansen. Towards Ontology Matching Suitable for Information Integration in Time-Critical Situations. DTIC Document, 2011.
- [2] E. Vysniauskas, L. Nemuraiteand A. Sukys. A hybrid approach for relating OWL 2 ontologies and relational databases. *International Conference on Business Informatics Research*. Springer, 2010: 86-101.
- [3] P. Chhaya, K. Leeand K. Shin et al. Using D2RQ and Ontop to publish relational database as Linked Data. *Ubiquitous and Future Networks (ICUFN), 2016 Eighth International Conference on*. IEEE, 2016: 694-698.
- [4] D. Calvanese, B. Cogreland S. Komla-Ebri et al., "Ontop: answering SPARQL queries over relational databases," *Semantic Web*, vol. 8, no. 3, pp.471-487, 2017.
- [5] M. Paradies, W. Lehnerand C. Bornhövd. GRAPHITE: an extensible graph traversal framework for relational database management systems. *Proceedings of the 27th International Conference on Scientific and Statistical Database Management*. ACM, 2015: 29.
- [6] A. Jindal, S. Madden, M. Castellanosand M. Hsu. Graph analytics using vertica relational database. *Big Data (Big Data), 2015 IEEE International Conference on*. IEEE, 2015: 1191-1200.
- [7] K. Cerans, G. Barzdinsand G. Bumans et al., "Relational Database Semantic Re-Engineering Technology and Tools," *Baltic Journal of Modern Computing*, vol. 2, no. 3, pp.183, 2014.
- [8] R. Winkler, "SpiderMass: semantic database creation and tripartite metabolite identification strategy," *J MASS SPECTROM*, vol. 50, no. 3, pp.538-541, 2015.

- [9] J. F. Sowa. *Principles of semantic networks: Explorations in the representation of knowledge*. Morgan Kaufmann, 2014.
- [10] C. Vicknair, M. Maciasand Z. Zhao et al. A comparison of a graph database and a relational database: a data provenance perspective. *Proceedings of the 48th annual Southeast regional conference*. ACM, 2010: 42.
- [11] R. S. Xin, J. E. Gonzalez, M. J. Franklinand I. Stoica. Graphx: A resilient distributed graph system on spark. *First International Workshop on Graph Data Management Experiences and Systems*. ACM, 2013: 2.
- [12] Z. Khayyat, K. Awaraand A. Alonazi et al. Mizan: a system for dynamic load balancing in large-scale graph processing. *Proceedings of the 8th ACM European Conference on Computer Systems*. ACM, 2013: 169-182.
- [13] N. Bronson, Z. Amsdenand G. Cabrera et al. TAO: Facebook's Distributed Data Store for the Social Graph. *USENIX Annual Technical Conference*. 2013: 49-60.
- [14] S. S. Singapogu, K. Guptonand U. Schade, "The Role of Ontology in C2SIM," *21st ICCRTS*,vol.,2016.
- [15] R. Haberlin, P. C. Da Costaand K. B. Laskey. Probabilistic ontology architecture for a terrorist identification decision support system. DTIC Document, 2014.
- [16] R. Haberlin, P. C. Da Costaand K. B. Laskey. An Ontology for Hypothesis Management in the Maritime Domain. DTIC Document, 2011.
- [17] P. H. Deitz, J. R. Michaelis, B. E. Brayand M. A. Kolodny. The Missions & Means Framework (MMF) Ontology: Matching Military Assets to Mission Objectives. *ICCRTS Conference*. 2016.
- [18] J. Park, E. K. Leeand Q. Wang et al. Health-connect: An ontology-based model-driven information integration framework and its application to integrating clinical databases. *Information Reuse and Integration (IRI), 2012 IEEE 13th International Conference on*. IEEE, 2012: 393-400.
- [19] P. Mitra, N. F. Noyand A. R. Jaiswal. Omen: A probabilistic ontology mapping tool. *International Semantic Web Conference*. Springer, 2005: 537-547.
- [20] N. Choi, I. Songand H. Han, "A survey on ontology mapping," *ACM Sigmod Record*,vol. 35,no. 3, pp.34-41,2006.
- [21] A. Moro, A. Raganatoand R. Navigli, "Entity linking meets word sense disambiguation: a unified approach," *Transactions of the Association for Computational Linguistics*,vol. 2, pp.231-244,2014.
- [22] S. Cucerzan, "Large-scale named entity disambiguation based on Wikipedia data,"2007.
- [23] J. G Enríquez, V. Lee, M. Goto, F. J Domínguez-Mayoand M. J Escalona. Entity Identification Problem in Big and Open Data. *Proceedings of the 17th International Conference on Enterprise Information Systems-Volume 1*. SCITEPRESS-Science and Technology Publications, Lda, 2015: 404-408.
- [24] E. Gabilovichand S. Markovitch, "Wikipedia-Based Semantic Interpretation for Natural Language Processing," *arXiv preprint arXiv:1401.5697*,vol.,2014.
- [25] S. Li, Z. Yang, J. Xiuand C. Liu. Research on an intelligent semantic-based information extraction framework. *Communications and Information Technologies (ISCIT), 2016 16th International Symposium on*. IEEE, 2016: 676-680.

- [26] N. Kambhatla. Combining lexical, syntactic, and semantic features with maximum entropy models for extracting relations. *Proceedings of the ACL 2004 on Interactive poster and demonstration sessions*. Association for Computational Linguistics, 2004: 22.
- [27] H. Jaeger, "Artificial intelligence: Deep neural reasoning," *NATURE*, vol. 538, no. 7626, pp.467-468, 2016.
- [28] A. Hunter, "A probabilistic approach to modelling uncertain logical arguments," *INT J APPROX REASON*, vol. 54, no. 1, pp.47-81, 2013.
- [29] E. Davis and G. Marcus, "Commonsense reasoning and commonsense knowledge in artificial intelligence," *COMMUN ACM*, vol. 58, no. 9, pp.92-103, 2015.
- [30] R. S. Michalski, J. G. Carbonell and T. M. Mitchell. *Machine learning: An artificial intelligence approach*. Springer Science & Business Media, 2013.
- [31] Q. Ding. Research on Data Extraction and Distributed Graph Data Management. Yunnan University, 2016.
- [32] A. DeFrancesco and B. McQueary. A Semantics-Based Approach to Schema Matching and Transformation in Network Centric Environments. *15th ICCRTS*. 2010.
- [33] Multilateral Interoperability Programme (MIP)[EB/OL][2017-09-15]. <https://public.mip-interop.org/Pages/Default.aspx>.
- [34] J. Webber. A programmatic introduction to neo4j. *Proceedings of the 3rd annual conference on Systems, programming, and applications: software for humanity*. ACM, 2012: 217-218.
- [34] J. Tang and Z. Luo, "Research on RDF data storage based on graph database Neo4j," *Information Technology*, vol., no. 6, pp.115-117, 2015.
- [35] S. Zheng, F. Wang and H. Bao et al., "Joint Extraction of Entities and Relations Based on a Novel Tagging Scheme," 2017.
- [36] Huang H, Li X, Wang X. "Key Issues in C-BML: A Survey," *Journal of Command and Control*, vol.1, no.3, pp.241-253, 2015.
- [37] Schade U, Biermann J, Frey M, et al. "From battle management language (BML) to automatic information fusion," *Information Fusion and Geographic Information Systems*, vol., no.10, pp.84-95, 2007.

9th KCSO & 22nd ICCRTS

“Frontiers of C2”

Where are the Knowledge Systems? Understanding obstacles to technology adoption

Track 10: Knowledge Systems for Coalition Operations
(Paper no: 76)

Jitu Patel and Peter Houghton

Defence Science and Technology Laboratory, UK

Dstl, Crown Copyright © 2017. This information is licensed under the terms of the Open Government Licence (<http://www.nationalarchives.gov.uk/doc/open-government-licence/version/3>).

For any inquiries regarding this paper, email jmpatel@dstl.gov.uk.

Abstract

Command and Control (C2) aspects of military operations, though very complex, knowledge intensive and usually time bounded, tend to be carried out with limited use of command and decision support tools. While there have been isolated successes such as the Dynamic Analysis and Replanning Tool used during operation Desert Shield, routine adoption of knowledge systems have remained low. The problem is not with the availability of technology, as defence (both Governments and Industries) have invested significant effort in tool development over the years. This paper will discuss findings from an earlier study, which investigated factors that may be contributing to the lack of exploitation of knowledge systems, and present insights from recent experience working in a standing joint force headquarters. The paper will conclude with an agenda for future research that contributes to our understanding of how best to improve technology adoption.

Introduction

Just over two decades ago, the ARPA/Rome Laboratory Planning Initiative (ARPI) organised annual Integrated Feasibility Demonstrations (IFDs) to foster transitioning of advanced technology, particularly knowledge systems¹, into operational systems (Tate, 1996). The first IFD was the Dynamic Analysis and Replanning Tool (DART), which was used during Operation Desert Shield by the US Transportation Command to support deployment planning (Bienkowski and Edwards 1996). DART is an example of a successful technology transition, but unfortunately successes such as this are transient and not universal.

If one were to walk into a division or brigade headquarters today, one would typically see a few large screen displays and rows of functionally-arranged tables with laptops. Each of the laptops would be networked, and some may be able to project information on the large screens. A more detailed look at the laptops would reveal that they have a plethora of applications, centred around standard office automation tools, but with a few others intended to support planning and decision making. However, a closer and continued inspection of which of the applications were being used would reveal that only the office automation tools (e.g., word processing, spreadsheet and drawing tools) are routinely used; and, one would be hard pressed to find much evidence of the use of anything else. Indeed, it would be more like to find that military staffs are stooping over paper maps, with some huddled next to whiteboards or flip-charts. This leads to an obvious question, why in this “information age” do people still prefer to carry out many of the primary activities manually?

One might have expected that within Command and Control (C2) aspects of military operations, given that they are very complex, knowledge intensive and usually time bounded would require more extensive use of command and decision support tools. The problem appears not to be with the availability of technology, as Defence (both Governments and Industries) has invested significant funds and effort in tool development over the years. The aim of this paper is to therefore to discuss findings from studies that have tried to identify those factors that may be contributing to the lack of exploitation of the tools, and also to present some insights from observing Joint Exercises over the years

¹ Planning and Decision Support Tools

and from more recent experience working in a standing joint force headquarters. The paper will conclude with an agenda for future research that contributes to our understanding of how best to improve take-up of technology in headquarters.

Challenges to Technology Adoption

In 2007 UK MOD commissioned a study on the challenges of technology insertion and the impacts that technologies have had upon organisations. One of the objectives of the study, relevant to the subject of this paper, was to review the state-of-the-art research on challenges to technology adoption. The study conducted a comprehensive review of the relevant literature and interviewed a number of subject matter experts. The literature review showed that the success rate of integrating technology with the business goals of an organisation was surprisingly low i.e. less than 50%. From the perspective of technology adoption, organisational factors such as the management support and the organisation's appetite for change appear to play a significant role in whether the technology will be used or not (Dawson 2007).

Ewusi-Mensah and Przasnyski (1991) argued that the success of a technology insertion project is strongly influenced by the degree of senior management involvement and the level of end-user participation. Related to the former, there is evidence that technology change succeeds when an organisation has both an appropriate political environment and change-favouring norms and culture (Tolbert and Zucker 1983). On the contrary, organisational politics can significantly hinder change programmes. For example, managers and staff can become major hurdles to a technology change if they believe their jobs are at stake, or they will be giving up some of their authority and control. To impede implementation processes people have used non-cooperative tactics such as "keeping out of the way", "withholding vital information" and demonstrating "non-availability" (Keen 1981).

Along with organisational culture and climate, Nielsen (1994) has argued that technology acceptability is also determined by users' perception of the usefulness of the technology which includes factors such as: how easy it is to learn, whether there are intuitive interfaces, the degree of memorability, quick recovery from errors, and overall user satisfaction. There is also an applicable theory: the Unified Theory of Acceptance and Use of Technology' (UTAUT), which aims to explain user intentions to use technology as well as explaining subsequent usage behaviour. The theory holds that four key constructs (performance expectancy, effort expectancy, social influence, and facilitating conditions) are direct determinants of usage intention and behaviour. Gender, age, experience, and 'voluntariness' of use are posited to mediate the impact of the four key constructs on usage intention and behaviour. (Venkatesh et al. 2003).

Factors Affecting Adoption

In 2014 Dstl commissioned a more focussed study to understand why technology utilisation, particularly of command and decision support tools, has remained at such a low level. The aim was to investigate factors that might be contributing to the lack of exploitation of the intended utility of command support tools. The study was conducted by a multidisciplinary

team, drawn from Industry and Academia that used a combination of case studies, stakeholder interviews, literature reviews, and existing trials data to generate and integrate insights (Jaya-Ratnam et al. 2014).

The data was collected from interviews, and subsequent analysis categorised barriers to tool utilization in terms of common themes and issues as shown in Table 1 below.

<i>Theme</i>	<i>Description</i>
REQUIREMENTS	Issues that indicate that User requirement for the tool (or function it is supporting) may not have been fully understood or realised during acquisition.
DESIGN	Issues with the tool related to poor consideration of human factors and usability in the software and/or hardware design, e.g. usability, error management, physical ergonomics, accessibility from workstation
PERSONNEL	Issues associated with User characteristics, noting that each tool can have a range of different users – operators, commanders, maintainers, etc.
IMPLEMENTATION	Issues that could have been avoided through improved implementation and management of change, e.g. timing, senior champions, understanding benefits
TRAINING	Issues with training and education requirements, course content, timing or availability, e.g. refresher training, access to theatre tools on exercises.
SUPPORTABILITY	Issues associated with the need for improved or increased technical support, e.g. database mangers, computer expertise or maintenance once in-service.

Table 1: Data categorisation themes²




Figure 1: Decision tool adoption, utilisation and utility "cascade"³

² Reprinted from Jaya-Ratnam et al. (2014 p. 10), with permission of DIEMconsulting Ltd.

Jaya-Ratnam et al. (2014) found that the identified themes could be viewed as steps in a cascade of development activities: Adoption – whether it is brought into “service”, Utilisation – whether users actually try to use it, and Utility – whether users gain a benefit, as shown in Figure 1 above. This cascade diagram has the potential to be a useful aid, in the form of prompting questions about necessary factors to consider along the development path of getting a tool successfully implemented and exploited. For example, the following are some issues that those attempting to acquire and implement a decision support tool must address:

- *Usefulness* - decision tools are more likely to be adopted if they carry out mundane and mechanistic aspects of a task, such as automatic capture of data/information, and leave the deeper sense-making and decision making to the human.
- *Access* – the tool is accessible as and when required.
- *Benefits and Dis-benefits* – users’ perception about the tool, beneficial or negative (i.e., dis-benefits). It was found that identified dis-benefits have greater negative impact on utility than meeting the requirements on benefit.
- *Training* – Users are adequately trained just before use (to ensure there is no skill fade). The assumption is that the tool is mature; otherwise, it would be ignored because of “dis-benefits” due to poor experience of bad design.
- *User Group* – There needs to be a critical mass of users and data to ensure there is material constantly being created and updated for exploitation. This should be supported by a User forum where users can go for help and support.

The study findings suggest that there is no clear and obvious relationship between benefit and adoption. However, there is a relationship between the level of negatives (i.e., dis-benefits) and non-adoption; adoption and utilisation was negatively correlated to level of “reported negatives”. This means that a current focus on equipment programmes of attempting to identify and verify requirements for improvement in decision making is not sufficient. It is also necessary to set requirements for any new system to not increase any “dis-benefits” for the user, which would then lead to non-adoption.

User Needs – Understanding the Requirements

From what we have observed over numerous years at UK Joint Exercises, and in particular from close observation of staff working at the UK’s Standing Joint Force Headquarters during 2016, staffs appear to prefer manual methods for information processing rather than exploiting the existing decision and planning support tools⁴. Whilst “ease of use” and “training” is partially the cause, they do not fully explain the reluctance to use the tools. We found in addition, the following two other factors that seem to influence persistent use of tools within a headquarters.

³ Reprinted from Jaya-Ratnam et al. (2014 p. 17), with permission of DIEMconsulting Ltd.

⁴ It appears that staff fall back on their experience of having no tools when they were Troop, Platoon and Battle Group Commanders.

Opportunity to Practise

One of the defining characteristics of a formation level headquarters is that there is a constant churn of staff every couple of years⁵. Furthermore, during Exercises, the headquarters numbers surge two to three fold due to augmentation⁶. While it is possible to ensure all staffs are trained before joining the headquarters, it is often the case that they do not get sufficient time to practise use of the tools intended to aid them. There are a number of reasons for this which includes:

- The battle rhythm - while the headquarters is in barracks its focus tends to be on administrative tasks and on refining deployment rather than on improving planning and decision making skills and competencies. It is quite normal to find that the primary opportunity for staffs to practise key HQ processes and use of tools are during exercise preparation and execution. However, during preparation stages the augmentees may not be present, and during execution, the core staff to augmentee ratio may be 1:2+. The end result is that the functional team leader tends towards the lowest common denominator and will choose not use any tool support (particularly if he/she is not an expert user and/or the tool is not intuitive).
- Lack of facility - even when time to practise is made available, the headquarters may not have a ready access to a facility to practise HQ processes using the tools (Patel and Patterson, 2017). Indeed, as noted above, the only time the headquarters has the opportunity to practise is during Exercises, which are focussed on assessing headquarters' competence and not on learning the tools of the trade. Hence, it is not uncommon to see staff conducting their work manually, as this is perceived to be by far the least risky option.

Scope and Functionality

Decision and Planning Tools are typically designed to support specific functions (e.g., TOPFAS⁷ in the case of military planning). Unfortunately, such tools have implicit and inbuilt assumptions on the ways of operating and these are frequently not easily adapted to new or different thinking and concepts. They also tend to impose a significant training burden and due to tool complexity and infrequent training cycles the staffs suffer from skills fade. Therefore, this factor again contributes to a lack of tool use.

It is frequently assumed within the Systems Procurement Community that functional areas within a headquarters can work independently. This is far from the truth. Headquarters, particularly the effective ones, provide a more coherent C2 capability where there is a more seamless dissemination of relevant information within the HQ. There are also robust and effective interfaces vertically (i.e., to higher and lower headquarters) and with external organisations. Thus, specialist functional tools that do not support a more integrated HQ

⁵ The duration may be shorter or longer for some staff. Typically the change would be at different periods during the year and either Service or rank based.

⁶ Augmentees may have the training and experience of tools, however this is not guaranteed. Also, what is quite common is that they would join the headquarters just before the start of an Exercise with little or no opportunity to practise the processes and/or tools with their new team members.

⁷ Tools for Operational Planning Functional Area Services (TOPFAS) is an integrated set of collaborative planning and decision support tools developed by NATO Communication Information Agency.

operation, and do not support coherent information dissemination between teams, will tend to be avoided. Instead, those that tend to be adopted are those that support generic functions such as information management, dissemination and sharing. However, even here, it is not given that such tools will be adopted and exploited, unless they are configured and used in a consistent manner across the headquarters. This frequently requires a skill sets that a military HQ does not have in abundance.

Planning and decision making is a collaborative activity involving staffs from across the various branches within an individual headquarters, as well as with other headquarters and external partner organisations. Staffs will therefore tend to favour those tools that are common across all of the potential collaborators. Thus, there is a tendency is to fall back, either completely on manual approaches (paper, maps, drawings, physical meetings and notebooks etc.) or on the lowest denominator technology such as office applications. Even these then tend to be used in a lowest common denominator manner i.e. with only a minimal level of functionality actually being used.

Conclusion

Over the last two decades a significant amount of research and development effort has gone into creation of knowledge systems (i.e., planning and decision support tools) in an attempt to improve the effectiveness of staff working in headquarters. Unfortunately, very few of these tools have been adopted. Getting the requirements right, and having a strong user involvement from the start, will both improve the chances of success. However, as noted above, there are many other factors that need to be attended to before tool adoption becomes more widespread. As noted previously in the successful example of tool adoption success:

“Transportation planners readily accepted DART because they had helped define the initial prototype capabilities, refine the prototype into the operations systems, and analyse elapsed planning and analysis times to quantitatively identify the major sources of improvement.” (Bienkowski and Edwards 1996, p 37).

Our more recent studies suggest that organisational and cultural factors also play an important role in technology adoption, and that this can be undermined by perception of dis-benefits. Observations of headquarters working suggest that many of the tools have been developed without clear understanding of the user needs and how the tools will be utilised within a headquarters. From the research perspective, there are two fundamental questions to be addressed:

- From a cost benefit and effectiveness perspective, would it be better to educate and help command staff continually practise new problem-solving and thinking skills instead of making large investments in technology-driven decision support tools?
- Should there be a different balance between investing in both education and tool development, where there is a shift of the balance away from technology via the use of much lighter weight, but more adaptable tools? If so, what would these tools look like and how would they be developed? Finally, what would the education and training to go with these tools be, and how would it be developed and sustained?

References

- Bienkowski M. and G. Edwards G. (1996). Demonstrating the Operational Feasibility of New Technologies: The ARPI IFDs in Tate, A. (Ed.) *Advanced Planning Technology: Technological Achievements of the APRA/Rome Laboratory Planning Initiative*, AAI Press, Menlo Park, California, 35-42.
- Dawson B. (2007). *The Impact of Technology Insertion on Organisations*, HFIDTC/2/12.2.1/1.
- Ewusi-Mensah, K. and Przasnyski, Z. (1994). Factors contributing to one abandonment of information systems development projects. *Journal of Information Technology*, 9, 185-201.
- Jaya-Ratnam D., Fletcher G., Smith T. & Schraagen J. (2014). *Final Report, TIN 3.092 C2 Concepts Task 1: Inadequate utility of command and decision support tools*, UC-DHCSTC_1386917_H_T3_092_1/002.
- Keen, P. (1981). Information Systems and Organisational Change, *Communications of the ACM*, 24, 24-33.
- Nielsen, J. (1994). *Usability Engineering*, Morgan Kaufmann Publishers.
- Patel J. & Patterson G. (2017). Headquarters Maturity Model: An approach to optimising a HQ for operational advantage. *Paper submitted to the 22nd International Command and Control Research and Technology Symposium*, Los Angeles.
- Tate, A. (Ed.) (1996) *Advanced Planning Technology: Technological Achievements of the APRA/Rome Laboratory Planning Initiative*, AAI Press, Menlo Park, California.
- Tolbert, P. and Zucker, L. (1983). 'Institutional Sources of Change in the Formal Structure of Organisations: The Diffusion of Civil Service Reform, 1880-1935'. *Administrative Science Quarterly*, 2, 22-39.
- Venkatesh, V., Morris, M.G., Davis, G.B., and Davis, F.D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27, 425-478.

2017 ICCRTS-KSCO

Title of Paper:

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

ICCRTS Topics:

Cognitive and Socio-technical Challenges
Interoperability, Integration and Security

KSCO Topics:

Coalition and Team Information Sharing
Coordination and Collaboration
Cultural Influences

Name of Authors:

David Mott
IBM UK
Southampton, UK

Troy D. Kelley
Human Research and Engineering Directorate
U.S. Army Research Laboratory
Aberdeen Proving Ground, MD, USA

Cheryl Giammanco
US Army Research Laboratory
Aberdeen Proving Ground, MD, USA

Soheil Eshghi
Department of Electrical Engineering and Institute for Network Science,
Yale University, CT, USA

Yunfeng Zhang
Human Agent Collaboration Group
IBM T. J. Watson Research Center, YorkTown Heights, NY, USA

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

Abstract

We describe research on understanding group mutability in the behaviour of external groups, and how interventions by coalition forces may affect the behaviour in terms of controlling hostile groups and encouraging friendly groups. We explore how emotion may influence the behaviour of individuals by affecting the type of reasoning that they undertake, encouraging "uncritical" rather than "critical" thinking. We describe a computational framework holding a cognitive model of an individual operating within a group context, inspired by theories from social science. Individuals relate to in-groups and out-groups and have beliefs that are associated with emotions. Cognitive Appraisal Theory is used to evaluate incoming memes "pronounced" by external speakers, appraising the effects of the memes on an individual's self-esteem taking account of their group relationships as indicated by social identity theory, and leading to an emotion in the individual. Appraisal is followed by a process of coping that seeks to handle the effects by either performing problem-focussed (critical) or emotion-focussed (uncritical) thinking, according to the current emotional state of the individual. This model is implemented within a Cognitive Architecture (Soar) as a set of reasoning processes that handle beliefs and emotion. The model is integrated into a multi-agent simulation tool (Repast Simphony) allowing the simulation of populations of individuals interacting and spreading rumours, or memes, together with interventions. We describe how this framework could be used to construct experiments to explore how different situations lead to group mutability and behaviour, together with the effects of interventions by coalition forces.

1 INTRODUCTION

One of the aims of the Distributed Analytics and Information Science International Technology Alliance (DAIS ITA) programme [DAIS-ITA 2016] is to understand the mutability of groups to support the achievement of desirable outcomes by the more accurate prediction of group behaviour and the design of effective intervention strategies. This paper outlines research to support the modelling of group mutability and factors causing behaviour, by exploring how emotion can affect the reasoning capabilities of group members. We seek to distinguish between "critical thinking" [Fisher 2011] where logical principles are employed, and "uncritical thinking", where these principles are not necessarily applied. Uncritical reasoning seems to underlie a range of undesirable phenomena such as conspiracy theories, demonisation of individuals and organisations, the refusal to consider expert opinion and the spreading of fear and false information, which in turn has the potential to adversely affect group behaviour. We explore the hypothesis that emotion can have a significant effect on cognition and on the use of uncritical thinking in particular; therefore the understanding of the effects of emotion on cognition is useful in understanding the mutability of group behaviour and how it might be influenced

In our research we are developing computational models based upon cognitive architectures and social theories, that simulate the effects of emotion on cognition in the context of group behaviour and thereby have the potential to provide characterizations of group mutability and the effects of intervention strategies. This work will also support other DAIS ITA research, including the development of a high level model of group behaviour and the development of a meta-heuristic framework for describing different modelling strategies.

The scope of this paper is to describe our computational framework for modelling emotion, cognition and group behaviour, with a focus on critical and uncritical reasoning. We provide an example of a basic experiment using this framework, and discuss how future experiments could be performed, but at this stage we do not provide detailed, scientifically validated, experimental results in scenarios involving complex social phenomena. The paper is structured in nine sections: section 2 reviews previous work in this area; section 3 describes a scenario used to focus the research; section 4 describes some basic concepts and associated social science theories used to develop the cognitive model; section 5 defines the cognitive framework and model in detail; section 6 shows how this framework integrates with a multi-agent simulation; section 7 describes a basic experiment using the framework; section 8 discusses extending the underlying cognitive architecture with emotion-specific modules; and section 9 draws conclusions and outlines future work.

2 PREVIOUS WORK

Previous work in computational models of emotions fall into four main categories: appraisal, dimensional, anatomical and rational [Marsella et al 2010]. The appraisal approach [Smith & Lazarus 1990] suggests that emotions arise from the cognitive agent's continuous assessments of the environment and the relationship between the environment and internal beliefs, behaviours, and concepts. The dimensional approach postulates that emotions are represented in a multi-dimensional space and vary along a continuum within the entire emotional space [Russell 2003]. The anatomical approach focuses on the neurological underpinnings of emotions, attempting to reconstruct neurological correlates and representations that are influenced by brain anatomy [Panskepp 1998]. The rational approach views emotions as another rational mechanism that aids cognition and decision making [Anderson & Lebiere 2003].

We aim for a cognitive approach, with an emphasis on how different forms of reasoning (uncritical and critical) may be influenced by emotion, yet accepting that there may be a logic to "emotional" reasoning. To this end, a combination of the appraisal, dimensional and rational approaches is being followed, rather than consideration of the neurological aspects of emotion.

As described below, our work combines a cognitive architecture (Soar), cognitive appraisal theory, and theories about social behaviour in a group context, with a view to understanding the effects of emotion on critical and uncritical reasoning. Here we survey previous research in the light of these topics. A number of researchers have integrated emotions and cognition within a cognitive architecture, with the dominant approach being based upon Cognitive Appraisal Theory [e.g. Smith & Lazarus 1990], where a situation is appraised using cognitive processes, leading to emotions, and is followed by coping strategies to deal with these emotions.

Much of this work is based upon the Soar cognitive architecture [Lehman et al 2006] to provide the computational framework in which to represent the appraisal process and its effects on reasoning. [Marinier et al 2009] describe an integration between cognitive appraisal theory and Soar, using appraisals such as goal relevance and suddenness to engender emotions and control the construction of actions in a game simulation. However the reasoning involved is aimed at a correct execution of the task; there is no attempt to model "uncritical reasoning" where emotion may "derail" the reasoning. In addition, the appraisals are defined by factors relevant to the game simulation, rather than by social theories about human behaviour in a group context. [Marsella & Gratch 2009] describe EMA, also based upon the integration of Soar and cognitive appraisal, focusing on the separation of (slower) cognitive process to represent the situation and (faster) appraisal processes. The coping mechanisms include uncritical reasoning, such as "mental disengagement" (as in coming to care less about a goal that has a low probability of success) as well as critical reasoning, such as "seeking information". However, the appraisal mechanisms are not derived upon theories about human behaviour in a group context; for example there is no coping mechanism to avoid damage to self esteem by ignoring information. [Laird 2008] describes an "Appraisal Detector" in Soar that implements cognitive appraisal theory as a sub-symbolic extension to the core Soar system, potentially allowing all cognitive functions, such as memory, to be affected by emotion, but in practice only supporting the learning function.

Other cognitive architectures have been used to model emotions. [Fum & Stocco 2004] uses ideas from ACT-R [Anderson et al 2004] to model the effects of emotion on human performance in a specific task, based upon the changing of the ACT-R memory activation mechanism to include emotional aspects. However the calculation of the emotional aspects is specific to the domain of the task and is not a generic mechanism based on a theory such as cognitive appraisal; nor does the research concern itself to the differences between critical and uncritical reasoning. [Lin et al 2011] describes EmoCog, a proposed cognitive architecture that includes an emotion component based upon appraisal theory, where memory nodes have associated emotional values that can focus attention (and hence cognitive processing) onto the most salient events and information. This work does discuss the

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

effects of emotion on the fidelity of the reasoning when executing a plan, where logically necessary steps are omitted due to low emotional arousal, but no specific appraisal mechanisms are given.

A different approach to the modelling of cognitive agents is that of BDI (Beliefs, Desires & Intentions), e.g. [Rao & Georgeff, 1995], which focuses on committing to goals in response to beliefs about the world and internal desires, and then detailed planning, and executing actions, to achieve these goals (the combination of commitment and plan being an "intention"). At this level of description, our work is similar to BDI in that we model agent behaviour that maintains beliefs about the world, has underlying desires and performs actions that are derived from satisfying these desires. However the BDI approach is generic and does not, in itself, provide a mechanism for determining desires (and the acceptance of beliefs) from social science theories of appraising and coping in a group context. Furthermore it seems to be taken for granted that BDI is concerned with modelling cognition that aims to correctly achieve specific tasks, rather than simulating the effects of emotion on uncritical reasoning.

The work surveyed above addresses many of our key issues, but none completely address our goal of simulating human behaviour which may consist of reactive responses to situations and may be based upon uncritical reasoning, false assumptions and cognitive biases. Our goal is to faithfully replicate and explain such behaviour, be it good or bad, in terms of the effects of emotion, rather than seeking to create effective problem-solving behaviour for specific tasks.

One body of research that addresses all of our target issues is that of Silverman, for example [Nye & Silverman 2013] describes the application of their cognitive architecture, PMFServ, to social learning. This contains an attentional mechanism based on social and other cues such as authority and influence of in-groups, and a motivation system for actions that is based on cognitive appraisal theory. Potential actions are appraised (via "activations") against a tree of goals, standards and preferences, some of which are social, such as "esteem", "treatment of out-groups" and "desirable future for the group" respectively. The activations are analysed to form a set of emotions which are then combined to form a subjective utility for each action, taking account of the expected change in emotions caused by an action. The action with the highest utility is then chosen for execution (though there are some other constraints that may rule out the choice of an action). Whereas this work addresses our target issues, at a more detailed level there are differences. Social group aspects of self-esteem and in- and out-group membership make up part of their appraisal mechanism, but these are linked in an indirect way via the mathematical formula for subjective utility, and there seems to be no specific symbolic chain of reasoning based upon a social theory such as Social Identity Theory to act as an explicit causal link between group information through the appraisal process to the emotional coping process. This may make it difficult to represent alternative reasoning mechanisms in the individual, such as the change from uncritical (emotional) to critical thinking. More generally, we are researching into the passing of rationale as part of the communication between individuals (as described in the section on trustworthiness below) and this requires explicit representation of the reasoning between inputs, appraisal, making of assumptions, and the resulting behaviour in terms of the memes to be communicated.

3 SCENARIO

A simple initial vignette (i.e. a detailed part of a larger contextual scenario) has been chosen that provides sufficient detail of an environment, individuals and groups, and has the capability of demonstrating mutable group behaviour together with opportunities for monitoring and intervention. This is a key concern of the coalition commander conducting civil-military operations to engage in, influence, or exploit relations between military forces, indigenous populations, and civilian organizations in support of stability and counterinsurgency within a host nation or region [Headquarters, Department of the Army 2013, Headquarters, Department of the Army 2014]. Coalition forces build trust and influence group behaviours by immersing themselves in the local

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

culture and politics, engaging with local leaders and the general populace to develop partner capacity-building programs focused on host nation governance and economic development [Department of the Army 2014]. Soldiers require cross-cultural language skills and an understanding of socio-cultural relationships in order to provide information messages that are culturally acceptable to the local public [Headquarters, Department of the Army 2013].

The initial vignette is based upon the passing of "pronouncements", or rumours, between individuals belonging to different, competing, groups; this allows the study of how "memes" [Dennett 1995] are taken up, spread or rejected by the communities, and how this affects reasoning both critical and uncritical. To provide some context and linking into the overarching scenario, a "back story" has been invented that provides some motivation for the individuals and the competing groups, based upon an ancient conflict between the "Reds" group and the "Greens" group in respect of land rights.

4 BASIC CONCEPTUAL MODEL

We define a conceptual model of the individual, including their set of beliefs, and an emotional "vector" that defines the individual's overall emotional state and the emotional strength of their beliefs. Groups are modelled with their antagonistic and collaborative relations to other groups, and the relationship of individuals to groups are defined in terms of the in-groups to which they identify and the out-groups with which they "un-identify". Communication between individuals is modelled as "pronouncements", where a speaker tries to pass on some information (a meme) to another individual in a face to face meeting.

Beliefs and pronouncements have linguistic semantic content in terms of a simple subject-action-object triple, and this content may be consistent with, or conflict with, the semantic content of other beliefs; for example "eats chocolate" and "bans chocolate" are inconsistent. Simple semantic reasoning is used to determine such consistency. We also define the emotional "content" of words which may be different for different groups.

The cognitive model takes theories from social psychology as a starting point, as outlined below. However a key part of the research is to develop some of the details in order to define a model that is computational, and our interpretation of these theories, as described in section 6, is relevant towards that aim. In some cases, intuitive hypotheses as to possible social effects have been devised. We do not suggest that the cognitive model presented here is complete and fully validated, but we do aim to demonstrate that the techniques described offer a means of representing and exploring different theories.

4.1 Cognitive Appraisal Theory

To construct a cognitive emotional model of the individual, we appeal to "Cognitive Appraisal Theory" [Smith & Lazarus 1990], which proposes how emotion and cognition are interconnected in two stages: appraisal and coping. In the appraisal stage, the current situation (i.e. the individual's relationship to the environment) is characterised across a number of "Appraisal Variables" including relevance and desirability in respect of goals, possible damage to self-esteem, unexpectedness, causal attribution, controllability, and how well the individual can cope with the situation. The values of the set of appraisal variables can then be mapped into a specification of an emotion. Cognition is involved in appraisal, since the world needs to be compared to the individual's beliefs, and inferences about complex factors such as causality are required. In the coping stage, the individual seeks to reduce the negative "damaging" effects of the appraisal; coping may be problem-focused, such as the use of negotiation, planning and logical problem solving; alternatively coping may be emotion-focused, such as the use of denial, shifting of blame or reducing the importance of the damage. Cognition is also involved in coping, since inference may be needed for solving problems and strategies may be needed for denial, etc. These two types of coping, problem-focused and emotion-focused, may be seen as applying critical thinking and uncritical thinking respectively.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

4.2 Social Identity Theory

However, Cognitive Appraisal Theory does not state how appraisal is affected by group membership, and it is necessary to employ further social science theories about the relationships between individuals and groups in order to build a more accurate cognitive model. For example, Social Identity Theory [Tajfel & Turner 1979] defines how individuals create a "social identity" which is part of their concept of "self" that is based upon the groups with they identify and un-identify¹. In this way, it is possible to model group dynamics, in the way that groups are formed and unformed.

5 COGNITIVE ARCHITECTURE FOR SIMULATING COGNITION

5.1 Overview

Since theory suggests that emotion and cognition affect each other, our computational model of cognition in the group context must take this into account. We use a Cognitive Architecture (CA), since this provides a predefined computational framework based upon theories of human cognition, but this must be extended to cover group-individual interactions and how they are affected by emotion. There are several candidate CAs, including ACT-R [Anderson et al 2004] and Soar [Laird 2008], but after review, Soar was chosen. In addition, there is existing research on the integration of emotion to cognition based upon Soar [Marinier & Laird 2007, Gratch & Marsella 2004, Marsella & Gratch 2009], although this does not address the effects on group behaviour.

Our work combines Cognitive Appraisal Theory and Social Identity Theory (including Self-categorisation Theory) to address emotion and cognition in relation to groups. An initial cognitive model has been constructed in the form of Soar rules, based upon the concepts defined in the conceptual model above together with an "implementation" of Cognitive Appraisal Theory and Social Identity theory, where appraisal is principally based upon self-esteem in relation to in- and out-groups, and coping strategies are based upon the limitation of damage to self-esteem. For example, a pronouncement is damaging to self-esteem when it is about an in-group but is inconsistent with the individual's beliefs. Coping with the damage may be undertaken by uncritical thinking emotional strategies such as rejecting the pronouncement or by critical thinking problem-solving strategies such as removing premises (e.g. being a member of an in-group) that lead to damaging appraisals. Choice of strategies is dependent upon the emotional fear level of the individual, high fear leading to uncritical reasoning and low fear level leading to critical thinking. The individual fear level itself is modelled as to be affected by the interactions between individuals, more interactions with out-group members leads to a higher fear level. It is calculated by the appraisal of the potential for physical (rather than emotional) damage to the self.

The cognitive reasoning performs appraisal of each contact with an external speaker in several ways. Firstly the pronouncement itself is appraised by comparing the semantics of the pronouncement with the individual's beliefs, using a simple semantic comparison, leading to an assessment of potential "damage" to self-esteem if it were to be accepted as true. Secondly the pronouncement is appraised to determine if potential "damage" to the physical self might occur, taking account of factors such as the

¹ More detail is provided in the related theory of Self Categorisation Theory [Turner et al 1987] which describes the cognitive process whereby individuals place themselves into categories (i.e. groups). An individual performs such categorisation by accentuating perceived similarities between members of the same category and perceived differences between members of different categories, using dimensions that the individual considers to be correlated to the categorisation. Furthermore, it explains the difference between an individual's "social identity" and a "personal" identity by the level at which the individual is self-categorising, a social identity being generated by a categorisation at a group level. Currently we do not implement a mechanism for accentuating the difference between groups, so do not apply Self Categorisation theory.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

emotional state of the speaker². After acceptance or rejection of the pronouncement, reasoning occurs to decide what pronouncement to pass on in the next contact with another individual; this is determined as the accepted belief with the highest emotional value.

The Soar model is described in more detail below, although the exact representation of the Soar rules and facts is not shown; rather the basic logic of the reasoning is described in informal terms.

5.2 Overview of Components

The Soar cognitive model, when performing uncritical thinking, is diagrammed in Figure 1, where rounded rectangles represent data, square rectangles represent reasoning processes and the arrows show how information flows between them:




Figure 1 The cognitive model in Soar, for uncritical thinking

The main components of the model are:

- Input/Output of pronouncements
- Linguistic translation and background information
- Linguistic synonyms based upon emotion
- Appraisal and coping of the pronouncement
- Choosing "My" pronouncement
- Appraisal and coping of the contact itself
- Maintenance of "My" state

The additional components involved in critical thinking are highlighted in the centre of Figure 2:

² We also model the emotional content of the words in the pronouncement, allowing the possibility that stronger emotive words could lead to a greater potential damage to the physical self.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

input and output pronouncement memes that convert these words to internal identifiers. In the current model this translation is limited, mapping those words that name individuals and groups (such as "Me", "Reds") to identifiers whilst leaving the other words "as is". However it opens up the possibility of more complex (and potentially group-based) linguistic processing.

5.4 Linguistic semantic relations

Two different memes may be analysed to determine if they are consistent or inconsistent with each other by simple analysis of the semantic content of the subject, act and object, and based upon basic semantic relationships between the concepts. Two types of inconsistency are defined. The first is "predicate inconsistency" where the act-object components define inconsistent acts³, for example

- the act "bans" is defined as being inconsistent with the act "eats", therefore the combination "bans chocolate" is inconsistent with the combination "eats chocolate". Thus the meme "Reds bans chocolate" is predicate inconsistent with the meme "Me eats chocolate"
- the act "hugs" is defined as being consistent with the act "loves", therefore "hugs kittens" is consistent with "loves kittens", and the meme "Greens hugs kittens" is predicate consistent with the meme "Me loves kittens"

The second type of inconsistency is "total inconsistency" where the complete subject-act-object components define inconsistent situations. For example:

- the meme "Reds bans chocolate" is total inconsistent with "Reds eats chocolate" (whereas the meme "Reds bans chocolate" is only predicate inconsistent with "Me eats chocolate", not total inconsistent)

5.5 Beliefs

Groups and individuals have sets of "beliefs", which are defined as:

- a meme, such as "Me eats chocolate"
- the agent (individual or group) that believes the meme (this may not necessarily be the same as the subject of the belief, thus "Me" can believe that "Reds ban chocolate").
- an emotion, being the emotional content of the meme with respect to the agent believer. Thus Me might believe that "Me eats chocolate" with a high degree of happiness (see below)

5.6 Emotion

The cognitive model currently represents six emotions⁴ [Ekman, 1992]:

- Happiness (or Joy)
- Fear
- Anger
- Sadness
- Surprise
- Disgust

These are represented in an "emotional vector" which is the ordered set of values for the emotions, each value being in the range 0 (no emotion) to 100 (full emotion). It is debatable whether all possible

³ The term "predicate" here is taken from linguistic practice where the verb-object combination is considered to be a predicate on the subject; thus the predicate "bans_chocolate(X)" is inconsistent with the predicate "eats_chocolate(X)".

⁴ [Plutchik, 1980] describes a further two, Trust and Anticipation

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

combinations are valid, but such constraints on the possible sets of values are not enforced. However, at this stage only Happiness and Fear are calculated and are treated as if they were opposites⁵, by combining them into a single intermediate value, a "Fear/Happiness level" which ranges from -100 (representing full happiness) to 100 (representing full fear); for example Fear/Happiness of -60 is represented as Fear=0, Happiness=60, and Fear/Happiness of 25 is represented as Fear = 25, Happiness = 0, and the Fear and Happiness levels are never positive at the same time. This intermediate combined value is only used to simplify the implementation of the Soar-based computations described below, and there is no logical requirement to manage multiple emotions as continua; indeed the model actually stores fear and happiness as two separate values.

Emotional vectors may be associated with different objects in the model, when that object is considered to have an emotional content. In the current model two types of object have associated emotional vectors:

- a belief, each belief having its own emotional vector
- "Me", representing the overall emotional state of the person being modelled.

An emotional vector has a "total emotional value", which is the summation of all emotional values in the emotional vector. (All values are zero or greater, so there is no offsetting of one emotion against another).

It should be noted that the approach using Cognitive Appraisal Theory does not hold emotions as being the fundamental unit of information; instead the emotions are derived from combinations of appraisal variables, see below.

5.7 Appraisal Variables and Vectors

Cognitive Appraisal Theory states that the external environment is first appraised in terms of a number of criteria that affect the individual and have the potential to raise emotions. Such criteria are called appraisal variables, and are each given a value. The following appraisal variables are currently being determined (together with their range of possible values)⁶:

- self-esteem [-100 for significant benefit, 100 for significant damage]
- physical safety [-100 for significant benefit, 100 for significant damage]

Within the cognitive model, the values of these variables make up an "appraisal vector", and such vectors may be associated with different objects and different aspects of these objects. Thus two appraisal vectors are calculated from each pronouncement, one is an appraisal of the content of the pronouncement (the meme), the other is an appraisal of the context itself within which the pronouncement was made.

The values of the variables contained in an appraisal vector are converted into an emotional vector, after applying the coping stage, which seeks to resolve any problems raised by the appraisal. The appraisal vector may also contain a "reason", that is a simple symbol that stands in for the cognitive reasoning (implemented as Soar rules) that led to this appraisal. This may be used in critical thinking to further review the logic of the appraisal and coping process, and to pass on reasons as part of the pronouncement, as described below. The vector may also contain additional information that is

⁵ This does not match the [Plutchik, 1980] circumplex of emotions where fear is opposed to anger and happiness is opposed to sadness. However in [Marsella & Gratch 2009] positive "desirability" is mapped onto hope and joy, whereas negative "desirability" is mapped onto fear and distress, or anger and guilt if causal attribution is involved. In the group context, we take self-esteem as a "desirable" situation, have followed a similar approach.

⁶ Cognitive Appraisal Theory defines more variables, and our future work may add to this list; one interesting (but complex) addition would be causal attribution.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

relevant to the type of appraisal. For example, the self-esteem appraisal based upon group membership (described below) will include the relevant group.

5.8 Trustworthiness

When critical thinking, the model aims to assess the trustworthiness of the source (i.e. an agent) of a meme, in order to help decide if it is to be believed. It is taken that memes from a trustworthy source will be believable. An agent receives memes via a pronouncement, and this provides two pieces of information that can assist the determination of trustworthiness of the source, the speaker and the reasons given for the meme. Since memes are passed between agents, we model two ways of how a meme is sourced: via an original trusted source; via another speaker who cites the trusted source as the reason for the meme.

To determine that a speaker is a directly trusted source, the current model considers that anyone who is "neutral" (i.e. does not belong to either the Greens or the Reds) is trusted⁷. To determine that there is an indirect trusted source for a meme, the reasons that are passed with memes are used. Thus a reason for a belief may include the fact that it was stated by a trusted source, hence the speaker is being a proxy for the original source.

Thus there are several stages in the propagation of trusted memes. In the first stage a directly trusted source (e.g. a neutral agent N) pronounces a meme M to another person P1, if the meme M is accepted by P1 the corresponding belief BM will have the trusted source N as a reason. In the second stage, P1 pronounces the meme M to another person P2 with the reason that P1 is proxying for the trusted source N. If P2 accepts the meme M then the corresponding belief BM will have the proxied source N as reason. Thus the meme is spread through those agents who are in critical thinking mode.

5.9 Logic of the Cognitive Model

We describe the cognitive model in more detail, based in part on the social science theories that underpin it (emphasised in bold), and in part on intuitive hypotheses as to the nature of cognition.

5.9.1 Inputs and outputs of pronouncements

The inputs and outputs to the model are pronouncements. On contact with another individual, the other individual's pronouncement is received, including the meme containing the semantic content, the identity of the speaker and the emotion expressed by the speaker. As a result of reasoning, "My" pronouncement (the one that "I" most wish to express) is output, including the meme and the emotion of the belief⁸. As described above, there is a linguistic translation between the pronouncement meme and an internal Soar representation.

5.9.2 Linguistic Semantics and Emotional Content of Words

Linguistic semantics (i.e. the semantic relationships between words as described above) and background information is used to compare the meme in a pronouncement to memes contained in

⁷ It is desirable to extend this to the modelling of authoritative, scientific reputation, although there are issues of how any authority is to be accepted by others. In addition, being a member of the same group does not of itself guarantee trust.

⁸ This raises the question as to the difference between the emotion of an individual (as per the input pronouncement) and the emotion of a belief (as per the output pronouncement). We believe that the individual has an overall emotion, which is separate from (but perhaps related to) the emotions of the beliefs; this would allow the representation of "undirected anxiety".

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

"my" beliefs, to determine which beliefs are consistent with the pronouncement and which are inconsistent (predicate or total) with the pronouncement⁹.

In addition, "emotional synonyms" are defined, that relate synonyms that have an associated emotional vector to basic words that do not. For example, the basic word "dislikes" has no emotional content, whereas the synonym "bans" has an emotional vector (containing fear). Emotional synonyms can potentially be used to estimate the emotion associated with the meme in a pronouncement, specifically by determining the emotional content of the "act" of the meme. Thus a meme such as "Reds bans chocolate" has a high emotional content, whereas "Reds dislikes chocolate" has a low emotional content.

5.9.3 Applying Cognitive Appraisal Theory to the pronouncement meme

The semantic content (meme) of the pronouncement, after determining its consistency or inconsistency with "my" beliefs, is examined **by applying Cognitive Appraisal Theory in its two stages of appraisal and coping**. Appraisal occurs to create the "self-esteem" appraisal variable, by assessing whether the action of a group (in or out) is consistent or inconsistent with "my" beliefs, **using Social Identity Theory**, that specifies that the concept of self is in part determined by the norms of the in-groups with which "I" identify and the norms of the out-groups with which "I" un-identify. This is interpreted in the following rules:

1. If "my" in-group performs an act that is (predicate) inconsistent with my beliefs then "my" self-esteem is damaged
2. If "my" in-group performs an act that is (predicate) consistent with my beliefs then "my" self-esteem is benefited
3. If "my" out-group performs an act that is (predicate) inconsistent with my beliefs then "my" self-esteem is benefited
4. If "my" out-group performs an act that is (predicate) consistent with my beliefs then "my" self-esteem is damaged

This assessment is determined by checking the subject of the pronouncement (to find the group performing the action) and by checking whether the act and object is predicate inconsistent with "my" beliefs, and results in a value for the "self-esteem" appraisal variable. This variable is associated with the meme in the pronouncement, rather than the overall emotional state of "Me". In effect the self-esteem damage is that which would occur if the pronouncement were to be accepted.

The logical reasoning for an appraisal using rule 1, including Social Identity Theory, is shown in Figure 3:

⁹ Currently we do not compare beliefs against each other, presuming that any set of prior beliefs given to the model are already consistent with each other.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning




Figure 3 Rationale for an appraisal of high self esteem damage

Here the speaker, who is a "Greens" states that the Greens ban chocolate. Since the speaker is trusted (because they are a Greens, see below for further discussion on this), it is taken that the Greens do indeed ban chocolate. In contrast "Me", who is also a Greens believes chocolate is likeable, which is predicate inconsistent with the Greens belief. Thus "my" belief is inconsistent with that of "my" in-group, Greens. Social Identity Theory states that one should follow the social norms, including the beliefs, of one's in-groups, but this is not happening, so self-esteem is damaged.

5.9.3.1 Coping and Critical Reasoning

Coping is then performed on the self-esteem damage appraisal. The cognitive appraisal theory suggests that there are two mechanisms for coping, emotion-based and problem solving-based; in the terms of this paper this corresponds to uncritical and critical reasoning respectively. In the model, two factors determine whether the agent is to perform critical reasoning or uncritical reasoning. The first factor is the current emotion vector of "Me", so that critical reasoning only occurs when the agent is at a lower level of fear. The second factor is the influence of the speaker's use of critical reasoning on

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

the listener, which we might call "Critical Reasoning Encouragement"¹⁰, where we hypothesise that an agent might be encouraged to perform critical reasoning (even tolerating a higher level of fear) by the knowledge that the speaker is itself performing critical reasoning.

Both of these factors are implemented in the following rules to determine if the listening agent is to perform critical or uncritical reasoning (and hence problem-solving or emotional coping)

- if the fear level is very high¹¹ then uncritical reasoning is performed
- if the fear level is high and the speaker is not performing critical reasoning then uncritical reasoning is performed by the listener
- if the fear level is high and the speaker is performing critical reasoning then critical reasoning is performed by the listener
- if the fear level is not high then critical reasoning is performed

Whether the speaker is, or is not doing, critical reasoning may be determined by the cognitive model of the listener by examining whether there are any reasons passed across with the pronouncement. Currently it is assumed that if reasons have been given for the pronouncement then the pronouncement has been arrived at by critical thinking, although this is a logical simplification.

Both types of coping must decide whether to accept or reject the pronouncement and (if accepted) to determine the resulting emotional vector for the belief. (A minor complication is that if the belief is already held by "Me" then the belief is not re-added nor is the existing belief's emotional vector updated). If the belief is accepted then the pronouncement is associated as the "source" of the belief.

In emotional (uncritical) coping, the pronouncement is rejected outright if the self-esteem damage is greater than 0; in this way the damage to the self-esteem is not taken (which would have been if the pronouncement were accepted). Pronouncements are accepted which are appraised with negative self-esteem damage, that way self-esteem is enhanced. In addition, the resulting emotional vector associated with the belief is determined from the self-esteem appraisal variable as follows:

- if self-esteem > 0 then this represents damage and the belief Fear/Happiness = damage
- if self-esteem <= 0 then this represents benefit and the belief Fear/Happiness is = -(damage)

Thus emotion of a specific belief arises only indirectly, calculated from the primary information in appraisal variables.

In some situations, coping also leads to a change in the overall person's emotional state as well as to a change in the emotion of specific beliefs. For example, as described in section 5.9.5, rejection of a damaging belief has a damaging "personal cost" in the calculation of the person's emotional state.

In problem solving (critical) coping, inspection of the reasoning that led to the appraisal can suggest how self-esteem damage may be avoided. In the example above, the damage is evaluated from two premises: that the speaker is to be believed; that "Me" identifies with the in-group Greens. Thus, logically there are two possible ways to remove the damage, by disbelieving the speaker, or by leaving the in-group. Consider the action of leaving the in-group, shown as a crossing out of the relevant premise ("identifies with") in Figure 4:

¹⁰ This is similar (but opposite) to emotional contagion, described below, where the emotion of the speaker affects the emotion of the listener.

¹¹ The levels that define high and very high fear may be set by the user in the Repast simulation environment described below.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning




Figure 4 Rationale after removing Greens as an "in-group", thus removing self-esteem damage

This results in the removal of dependent inferences, including the appraisal of high damage (shown as lines crossing out the inferences). This critical thinking coping strategy may be expressed as:

- if the reason for the self-esteem appraisal is that the associated in-group is being stated as having an inconsistent belief and the speaker is trusted, then the group is removed as one of "my" in-groups, and the pronouncement is accepted

This contrasts with uncritical thinking coping, where the pronouncement is just rejected out of hand with no further analysis. Consideration of the diagram shows that this would lead to an inconsistent state of knowledge where a speaker is both believed (because the speaker is trusted) and not believed (because the pronouncement is rejected), but this presumably would not be noticed by the individual, otherwise they might engage in critical thinking to resolve the inconsistency.

Another logical possibility is to not believe the speaker, which would remove the inference that "Greens ban chocolate" leading to the removal of the appraisal from the diagram. This is currently implemented in a different way, by determining the trustworthiness of the speaker prior to making the appraisal, and by only performing the appraisal if the speaker is trustworthy. Therefore no specific coping strategy is needed, since the appraisal has already taken this into account. This shows that there are alternative mechanisms for implementing the logic of appraisal and coping.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

At this stage of the research, these logical considerations in support of the appraisal and critical coping strategies have been done "by hand" leading to the design of Soar rules in the model. In previous work, [Mott et al 2010, 2015] we analysed the rationale and automatically calculated the premises and assumptions that led to certain conclusions. But this requires the underlying reasoning engine to have capabilities for examining the rationale graph, and this would require modifications to the Soar inference system. Such modifications could lead the cognitive architecture towards the concepts of argumentation theory [Dung 1995]. Some initial work has been done in the recording of reasons for beliefs and the passing of reasons with the pronouncement itself, in order to model the use of reasons as attempting to convince others of the validity of statements, [Sperber and Mercier, 2017] thus being part of the cognitive and emotional processes.

5.9.4 Determining "My" pronouncement

After the pronouncement has been accepted or rejected, the current set of beliefs is examined to see which has the most emotional salience and therefore is the one that "I" choose to pass on to the contact in return. This choice¹² is made as follows:

- prefer a belief whose total emotional level is higher
- of beliefs whose total emotional level are the same, prefer a belief whose source is a pronouncement¹³

5.9.5 Applying Cognitive Appraisal Theory to determine the person's emotional state

Intuitively, the contents of the pronouncement and the nature of the contact could have an effect on the overall emotional state of the person, as well as on the damage to self-esteem described above. An additional appraisal variable is used for this purpose, that of the potential damage to the person's physical safety (called here "self-safety") as a result of the contact. Such an appraisal must be coped with and may lead to an effect on the person's overall emotional state. There are a number of factors involved in the contact and the pronouncement that could suggest potential damage to self safety as assessed by the individual, and this assessment could be affected by the emotional state of the individual. One factor relates to the emotional state of the speaker, on the grounds of "emotional contagion" where a listener may be "contaminated" by the speaker's emotion; this information is directly available from the pronouncement¹⁴. Another factor is the potential effect of the semantic content of the meme on the person's emotion,¹⁵ though this can only be derived from the assessment of the meme against the person's beliefs in some way. Various alternative sources of appraisal of self-safety have been considered, and the model is still being developed in this area. The current appraisal of "self-safety" uses the following factors based upon the emotions that can be derived from the pronouncement:

1. the emotional state of the speaker of the pronouncement; a greater level of fear suggests a greater potential of damage to self-safety, whilst a greater level of happiness suggests a reduced potential of damage to self-safety
2. the emotional value of a belief resulting from accepting a pronouncement; a greater level of fear suggests a greater potential of damage, whilst a greater level of happiness suggests a reduced potential of damage to self-safety

¹² It would be possible to change the model so that more than one belief is passed on, leading to different simulation results.

¹³ This approximates to preferring more recent beliefs; it may be better to maintain a "recency" value for beliefs

¹⁴ In theory it could also be assessed from the emotional content of the words in the pronouncement meme

¹⁵ If we do not model this connection, then the content of the memes being passed are totally irrelevant to the person's emotional state; not only is this somewhat implausible, it also disallows the possibility that interventions could ever be based upon the propagation of information.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

together with a factor derived from the coping process of accepting or rejecting beliefs that was described above, based upon the idea that there is a "personal cost" to rejecting damaging beliefs¹⁶:

3. the damage to self-esteem, if the rejected pronouncement meme had been accepted

The appraisal value for "self-safety" is thus a combination of these three factors. At the coping stage, if the person is in a state of uncritical thinking, the person's overall emotion is increased by a proportion of this self-safety appraisal value:

- the person emotion is calculated from the old Fear/Happiness level (prior to the contact) + $K * \text{the self-safety appraisal}$.

where K is somewhat arbitrarily set to 0.2. However, if the person is critical thinking, then the appraisal of potential damage to self-safety is ignored, and the person's emotion does not change.

5.9.6 Maintaining "My" state

Various pieces of information held within the model represent the "My" current state and these may change over time, as more contacts occur. Such changing information comprises:

- the groups with which "I" identify (in-groups) and un-identify (out-groups)
- "my" beliefs (including a meme and the emotional state of that belief)
- "my" overall emotional state
- the reasons for "my" beliefs, including trustworthiness

The current implementation of the Soar model does not maintain this information over time, but instead the state is saved after each contact via the Java-based wrapper, and the Soar system is re-initialised with this saved state prior to the next contact¹⁷.

6 INTEGRATION TO MULTI-AGENT SIMULATION

The running of the Soar cognitive model implements an "interaction" between two individuals when a pronouncement is made, leading to the acceptance or rejection of the pronouncement and the readying of the individual to pass on their own pronouncement at the next contact. We have integrated this individual behaviour into a larger scale simulation with multiple agents, to provide a framework for exploring how larger scale behaviour might "emerge" from the individual interactions and exchanging of pronouncements. For this purpose REPAST Symphony [North et al 2013] is used, which provides a platform for developing multiple agents operating in a spatial environment, for the running of simulations and for gathering and analysis of data. We have integrated the Soar Cognitive Architecture into REPAST Symphony, providing a number of separate agents each behaving according to the cognitive model and passing pronouncements to other agents, allowing the running of simulation experiments.

7 EXPERIMENTATION

A simple experiment using the Soar/Repast simulation can be used to demonstrate how the framework can be used to model emotional and uncritical reasoning, and how external monitoring and intervention strategies could be simulated, although this experiment is not intended to be complete and scientifically valid.

The Repast time series graph in Figure 5 shows two alternative "bans" memes ("Reds bans chocolate", "Greens bans chocolate") flowing through two populations of agents, in two different groups

¹⁶ An alternative approach is for the self-esteem coping to accept the damaging belief with a fear level, but to record it is being false so that it is not employed in any of the logical reasoning

¹⁷ This is an implementation detail; logically it should be possible to maintain the state within Soar.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

(approximately) 50% Green and 50% Red. These are initially seeded by two pronouncers, one Red and one Green. As per section 5.9.3, these "bans" memes are happy for the "opposite" group and fearful to the "same" group (e.g. "Reds bans chocolate" is happy for the Greens and fearful to the Reds). All agents start with two "like" memes that will be the default pronouncements if the other memes are not accepted ("Me eats chocolate" and "Me likes kittens") These "like" memes are slightly happy to all agents irrespective of the group to which they belong. The initial average fear level of the entire population is 50 (the maximum possible fear level being 100), and the fear threshold below which critical thinking occurs is set so that 67% of the agents start in critical thinking mode.

The graph shows nine lines each displaying the change over time of a particular variable. Numbered from 1 to 9, these lines represent: 1) the overall percentage of the agents that are pronouncing any meme at all, 2) the percentage of agents that are pronouncing the default "likes" meme (e.g. "Me eats chocolate"), 3) the percentage of agents undertaking critical thinking, 4) the percentage of agents belonging to the Red group, 5) the average fear level of all agents 6) the percentage of agents belonging to the Green group, 7) the percentage of agents that have just switched to a different pronouncement meme, 8) the percentage of agents pronouncing the "bans" meme "Reds bans chocolate", 9) the percentage of agents pronouncing the "bans" meme "Greens bans chocolate".




Figure 5 Time series of key variables in a basic Repast Symphony simulation experiment

Several phenomena are visible in this graph and we can relate these to aspects of the model described above. Analysis of the simulation data shows that there are fairly complex interactions between the events, so we only attempt to show some of the main features here.

Up to about tick count 2000, it is mostly the default "likes" memes (2) that are taken up, since there are only two initial pronouncers of "bans" memes, and the critical thinkers will reject the "bans" memes (they are not spoken by a trustworthy source), leading to most agents pronouncing their default meme. Critical thinking encouragement increases the percentage of critical thinkers (3), which ought to inhibit fear increasing. However the small uptake of the "bans" memes that does occur will tend to increase fear (for reasons described in the next paragraph) and this offsets the increase of critical thinkers, causing a small overall increase in fear (5).

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

From this time on until about tick count 12000, the reach of the "bans" memes increases significantly (8, 9), and thereby eats into the take up of the default "likes" memes (2). As more interactions between the agents occur, "bans" memes are accepted by more non-critical thinking agents (of the "opposite" group) and then pronounced to other agents; some will accept (further increasing the reach) but others (non-critical thinkers of the "same" group as the "bans") will reject the memes. This rejection will increase fear level due to the personal cost for rejecting damaging beliefs (section 5.9.5). Thus the spreading of the "bans" memes will actually cause an increase in fear, and this is seen in line (5). As fear goes up, critical thinking is reduced (3), and the process of critical thinking encouragement is stopped, thus releasing the possibility of "bans" meme acceptance and fear increase.

There is an interesting "micro event" at tick count 12000, where the final acceptance of a "bans" meme takes place (9). This leads to a blip of increased fear (5), and detailed analysis of the data (not shown) indicates that this event causes ripples of increased acceptance and rejection of a bans meme, with its associated increase in fear, over a number of ticks, eventually leading to quiescence. After this, from about tick count 14000 to the end of the run, the values of the main variables do not vary a great deal. No further changes to the percentage reach of the "bans" (8,9) and "likes" (2) memes takes place, and analysis shows that all critical thinking agents hold the "likes" meme and all non-critical thinking (i.e. more fearful) agents hold one or other of the "bans" memes.

These initial simulations are intended to test the framework rather than representing validated scientific results in specific psychological situations. Nevertheless some patterns are beginning to emerge, such as the cyclic swapping of memes, the opposition of positive and negative forces on the fear level and critical thinking levels noted above, and the removal of inconsistent beliefs in the population seeded by a trusted source (not shown). The framework offers a number of parameters that may be varied and which can lead to different group behaviours. Some parameters are contained in the cognitive model (such as the approach to handling self-esteem and self safety and the nature of semantic information about the memes), and other parameters are contained in the simulation itself (such as the make up of the various populations and the groups to which they belong and the memes that may be passed around).

We are extending our experiments to model situations where emergence of relevant and real social phenomena may be hypothesised. Comparing the results of these simulations with the results of real-world studies will serve as a method of validating the models, as well as showing the minimal conditions under which the social phenomena can be observed. For example, we are planning to test the emergence of "false consensus" and "pluralistic ignorance" in these settings. Pluralistic ignorance arises when most individuals in a population privately reject a norm, yet keep their rejection private and conform in public because they misconstrue the public conformity of others as an expression of their private belief. [Bicchieri 2005] explains how non-transparent communication is a condition for pluralistic ignorance to arise. This non-transparency may, for example, be due to fear or bias, which both impede the aggregation of correct information. The current cognitive model incorporates fear, communication and bias, and thus can be used as the starting point for minimal experiments to find test conditions and parameters under which pluralistic ignorance is observed. Relevant parameters include the amount of private belief in the "taboo" meme in the population, the overall fear of the population, the size and number of groups (which affects both fear and bias), and the starting level of critical reasoning in the population. The outcome to be observed is the public pronouncements of individuals: we are interested in cases where there is a significant discrepancy between the empirical distribution of these pronouncements and the distribution of private beliefs in the population, which would be evident of pluralistic ignorance. False consensus arises when an individual assumes that their beliefs are shared by others, e.g. by people in general or by a specific in-group [Ross et al 1976]. To model this, we are extending the model to include the making of assumptions by individuals (for example of what a group believes) and the challenging of such assumed beliefs in order to achieve desirable interventions [Berkowitz 2005].

We also aim to explore additional variables expressed in social psychology theories within the Soar models, such as dynamic social impact theory which provides an explanation for socially transmitted

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

beliefs and development of culture through communication that considers, in addition to the strength, immediacy or proximity, and number of sources exerting social influence, the self-organizing properties of groups, such as spatial clustering and correlation of attitudes, consolidation of minorities, and continuing diversity [Latané, 1996].

8 Extending the cognitive architecture with emotional algorithms

Our current approach is to build the uncritical reasoning aspects on top of the cognitive architecture, as in [Marsella & Gratch 2009], rather than building emotional mechanisms within the architecture itself as in [Marinier & Laird 2004]. However the latter approach has the advantage of modelling the effect of emotion over the system as a whole, involving modules such as memory and attention.

Other researchers have leveraged techniques to handle emotions within cognitive architectures, and it is possible to use these as inspiration for similar techniques within a group context rather than within models of individuals. For instance, [Pirolli 2005] used spreading activation algorithms, which specify the spread of activation strength from one memory to the next, as a model of information scent, and developed models which manipulated the cost-benefit, or utility of the agent moving from one piece of information to the next. [Reitter & Lebiere 2011] used ACT-R memory algorithms as a basis of language evolution, allowing the researchers to model the changes within languages and amongst different groups. [Reitter & Lebiere 2012] used memory decay as a model of information decay within groups, in order to simulate group decision making and social cognition. There are similarities between the base level learning algorithms within ACT-R and the happiness algorithms developed by [Rutledge et al 2014], both of which are time based exponential functions subject to some kind of decay. [Long, Kelley & Avery 2015] successfully expanded the Rutledge model to include additional emotions besides happiness - these additions included fear, anger, sadness, disgust, and surprise. These models allowed for stochastic behaviors to be exhibited by a robot while the robot was executing a navigation task. In order to implement such techniques, it may be necessary to extend the cognitive architecture itself.

9 CONCLUSIONS and FUTURE WORK

We have constructed an initial cognitive model for expressing aspects of group behaviour, based upon social science theories and running on a cognitive architecture, and we believe that the representation of the theories of Cognitive Appraisal and Self Identity combined within a computable cognitive model provides a powerful basis for the explanation and simulation of the effects of emotion on critical and uncritical reasoning.

However, the cognitive model covers only a few aspects of group dynamics and requires extension in a number of areas. Firstly, although the integration of Soar and Repast supports the effects of the group on the individual, in that the group information and group beliefs will be present in each individual simulation, it is also desirable to allow the individual to affect the group (and hence other individuals indirectly), in which case a separate model of the group as it changes over time will be necessary. Such a model could be constructed as a separate Cognitive Architecture with its own group level cognitive model, or could be constructed within Repast providing some group level properties that change over time; in either case a mechanism is necessary to pass the changing group-level information down to the individuals. Secondly, the appraisal could be extended to cover more appraisal variables, such as causal attribution and its effect on perceived damage to self-esteem and the blocking of individuals goals. Thirdly the connectivity of the network could be extended with specific patterns of connection, allowing the exploration of social motifs or geographic-based communications. Fourthly, additional properties and relations amongst groups and individuals could be modelled, including roles of an individual in relation to a group, the status of the group as perceived by the individual and the status of the individual within the group. Fifthly, mathematical calculations could be used to maintain the strength of beliefs based upon recency and degree of exposure, as well as to apply the activation spreading algorithms to influence attentional mechanisms.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

We have integrated the cognitive model and architecture into a multi-agent simulation system, allowing the construction of experiments, where parameters may be adjusted to cause different group behaviours. It is important that we construct more detailed and scientifically valid experiments using the principles described above, in order to test the validity of the model and its ability to support the prediction of the mutability of the groups and the effects of interventions. In the long term this research aims to support our understanding of how select information, such as information about health services, veterinary aid, and infrastructure development activities may be conveyed during Military Information Support Operations to favourably influence the local population's attitudes, emotions, and reasoning about civil-military operations and US policy [Headquarters, Department of the Army 2013], as well as understanding how cognitive biases may affect reasoning in the coalition as well as the external groups.

10 ACKNOWLEDGEMENTS

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copy-right notation hereon.

11 REFERENCES

- Anderson, J. R. & Lebiere, C. (2003) The Newell Test for a theory of cognition. *Behavioral and Brain Sciences*, 26, 587-640.
- Anderson, J. R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., & Qin, Y. (2004) An Integrated Theory of the Mind, *Psychological Review*, Vol 111(4) 1036-1060
- Berkowitz, A.D. (2005) An Overview of the Social Norms Approach, in Lederman, L.C. & Stewart, L.P (Eds.) *Changing the Culture of College Drinking: A socially Situated Health Communication Campaign*, Hampton Press Inc.
- Bicchieri, C. (2005). *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press.
- DAIS-ITA (International Technology Alliance in Distributed Analytics and Information Sciences) (2016), <https://dais-ita.org>
- Dennett, D.C. (1995) *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon and Schuster.
- Dung, P. M. (1995) "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games." *Artificial Intelligence*, 77: 321-357 (1995).
- Ekman, P (1992). "An Argument for Basic Emotions". *Cognition and Emotion*, 1992, 6 (3/4), 169-200
- Fisher, A., (2011) *Critical Thinking: an Introduction*, Cambridge University Press
- Fum, D., & Stocco, (2004) A. Memory, Emotion, and Rationality: An ACT-R interpretation for Gambling Task results, *Proceedings of the Sixth International Conference on Cognitive Modeling*, 172-177, Mahwah, NJ: Lawrence Erlbaum
- Gratch, J., & Marsella, S. (2004) A domain-independent framework for modeling emotion. *Cognitive Systems Research*, 5(4), 269-306.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

Headquarters, Department of the Army (2013). Civil-Military Operations (Joint Publication 3-57). www.dtic.mil/doctrine/new_pubs

Headquarters, Department of the Army (2014). *U.S. Army Functional Concept for Engagement* (TRADOC Pam 525-8-5). <http://www.tradoc.army.mil/tpubs/>

Laird, J. E., (2008) Extending the Soar Cognitive Architecture, *Proceedings of the 2008 conference on Artificial General Intelligence*: Memphis: IOS press

Latané, B. (1996) Dynamic Social Impact: The Creation of Culture by Communication, *Journal of Communication*, 46(4), 13-25.

Lehman, J.F., Laird, J.E., Rosenbloom P.S., (2006) A gentle introduction to Soar, an Architecture for Human Cognition: 2006 update, *Invitation to cognitive science* 4, 212-249

Lin, J., Sparagen, M., Blythe, & Zyda, M. (2011) EmoCog: Computational Integration of Emotion and Cognitive Architecture, *Proceedings of the 24th International Florida Artificial Intelligence Research Society Conference*

Long, L. N., Kelley, T. D., & Avery, E. S. (2015). An Emotion and Temperament Model for Cognitive Mobile Robots. In *24th Conference on Behavior Representation in Modeling and Simulation (BRIMS)*, Washington, DC.

Marinier R.P. (2008) *A Computational Unification of Cognitive Control, Emotion, and Learning*, Dissertation, University of Michigan, Ann Arbor

Marinier, R.P., & Laird, J.E., (2004) Toward a Comprehensive Computational Model of Emotions and Feelings, *Proceedings of the Sixth International Conference on Cognitive Modeling*, 172-177, Mahwah, NJ: Lawrence Erlbaum

Marinier, R. P., & Laird, J. E. (2007) Computational Modeling of Mood and Feeling from Emotion. *Proceedings of 29th Meeting of the Cognitive Science Society*.461-466. Nashville: Cognitive Science Society.

Marinier, R.P., Laird, J.E. & Lewis, R.L. (2009). A computational unification of cognitive behavior and emotion. *Cognitive Systems Research* 10 (2009) 48-69

Marsella, S.C., & Gratch, J. (2009) EMA: A process model of appraisal dynamics, *Cognitive Systems Research*, 10, 70-90.

Marsella, S., Gratch, J. & Petta, P. (2010) Computational Models of Emotion. In Scherer, K. R., Banziger, T., and Roesch, E. (Eds.), *A Blueprint for Affective Computing: A sourcebook and manual*, Oxford University Press.

Mott, D., Giammanco, C., Dorneich, M.C., Patel, J., & Braines, D. (2010). "Hybrid Rationale and Controlled Natural Language for Shared Understanding". Proceedings of the 6th Knowledge Systems for Coalition Operations, Vancouver, Canada.

Mott, D., Shemanski, D.R., Giammanco, C., Braines, D., (2015) Collaborative human-machine analysis using a Controlled Natural Language, SPIE Next-Generation Analyst III, April 2015. Also <http://nis-ita.org/science-library/paper/doc-2852>

North, M.J., Collier, N.T., Ozik, J., Tatara, E.R., Macal, C.M., Bragan, M. & Sydelko, P. (2013) Complex Adaptive Systems Modeling with Repast Symphony. *Complex Adaptive Systems Modeling* 2013, 1:3

Nye, B.D., & Silverman, B.G. (2013) Social learning and Adoption of New Behavior in a Virtual Agent Society, *Presence*, Vol. 22, No 2, Spring 2013, 110-140

Panskepp, J. (1998) *Affective Neuroscience: The Foundations of Human and Animal Emotions*, New York, Oxford University Press.

A Framework for Modelling the Effect of Emotion on Uncritical Reasoning

- Pirolli, P. (2005). Rational analyses of information foraging on the web. *Cognitive science*, 29(3), 343-373.
- Plutchik, Robert (1980). A General Psychoevolutionary Theory of Emotion, In Plutchik, R., Kellerman, H., (Eds.) *Emotion, Theory, research, and experience: Vol. 1. Theories of emotion*, Academic Press
- Rao, A. S., & Georgeff., M.P. (1995). BDI-agents: From Theory to Practice". Proceedings of the First International Conference on Multiagent Systems (ICMAS'95).
- Reitter, D., & Lebiere, C. (2011). How groups develop a specialized domain vocabulary: A cognitive multi-agent model. *Cognitive Systems Research* 12, 175–185.
- Reitter, D., & Lebiere, C. (2012). Social Cognition: Memory Decay and Adaptive Information Filtering for Robust Information Maintenance. In *AAAI*.
- Russell, J. A. (2003) Core affect and the psychological construction of emotion. *Psychological Review*, 110, 145-172.
- Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *Proceedings of the National Academy of Sciences*, 111(33), 12252-12257.
- Smith, C.A. & Lazarus, R. (1990) Emotion and Adaptation. In L. A. Pervin (Ed.), *Handbook of personality: Theory and Research* (pp. 609-637). New York: Guilford Press.
- Sperber, D. & Mercier, H. (2017) *The Enigma of Reason: A new Theory of Human Understanding*, Allen Lane.
- Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worchel. *The social psychology of intergroup relations*. Monterey, CA: Brooks/Cole. pp. 33–47.
- Turner, J. C., Hogg, M. A., Oakes, P. J., Reicher, S. D. & Wetherell, M. S. (1987). *Rediscovering the social group: A self-categorization theory*. Oxford: Blackwell

Privacy Technologies for Controlled Information Sharing in Coalition Operations

**Karen Myers^{*}, Tim Ellis^{*}, Tancrède Lepoint^{*}, Ronald A. Moore[^],
David Archer[#], Grit Denker^{*}, Steve Lu[†], Stephen Magill[#], Rafail Ostrovsky[†]**

^{*}SRI International
Menlo Park, CA
{<firstname>.<lastname>}@sri.com

[^]Pacific Science & Engineering Group
San Diego, CA
ronmoore@pacific-science.com

[#]Galois, Inc.
Portland, OR
{dwa,stephen}@galois.com

[†]Stealth Software Technologies, Inc.
Los Angeles, CA
{rafail,steve}@stealthsoftwareinc.com

Abstract

Information sharing among coalition partners must balance the benefits that can accrue from improved coordination with the risks of releasing information that ideally would be kept private. We consider how advanced privacy technologies can enable improved information sharing among coalition partners by both providing increased control over how information is used or released, and enabling principled characterizations of the impact of individual and cumulative sharing activities. We describe this work in the context of a humanitarian aid and disaster relief scenario, showing how the technologies can enable significantly increased and informed sharing.

1. Introduction

Information sharing is a major challenge for coalition operations. Coalitions can range in composition from single-nation, inter-service or inter-agency teams to large, multi-national groups augmented with non-governmental organizations (NGOs) and corporations. Members can range from close allies to infrequent collaborators to adversaries or competitors. Furthermore, these relationships can change abruptly, underscoring the need for flexibility and adaptiveness.

Effective coordination with partners can require the intentional release of information that ideally would be held private, given the anticipated benefits that can result. However, information security mechanisms developed for the military have been designed to impede rather than facilitate sharing, due to concerns over unintended consequences of information releases. The

cryptography community has made significant strides in recent years in developing advanced technologies that can be leveraged to safeguard privacy (for example, see the description in [Archer et al., 2016]). This paper describes an exploration into how these types of technologies can be employed to enable informed and controlled information sharing within coalitions.

To ground our work, we have been considering a use case rooted in humanitarian aid and disaster relief (HADR). We chose to focus on HADR for several reasons. First, it is representative of real-world, multi-nation coordination tasks that happen on a regular basis. Second, it encompasses privacy concerns at multiple levels: individuals, intra-organization, and inter-organization, with organizations spanning nation-based, commercial, and NGOs. Third, it supports a range of challenging privacy problems, including access to both structured and unstructured data, and multi-party coordination tasks that require sequences of information exchanges and joint computations.

At the heart of our approach is a platform called PRIME (Privacy-preserving Information Mediation for Enterprises). PRIME provides privacy management by integrating a set of privacy controls, comprised of security mechanisms and policy setting capabilities for data owners, as well as various analysis tools for measurement and prediction of information leakage. Beginning with requests from authorized users, PRIME manages request processing using permitted data and services to provide a response that satisfies the requester’s needs while remaining in compliance with the privacy requirements of data owners. Many tasks within coalitions involve ongoing, temporally extended coordination. For this reason, our approach adopts a *process-oriented* perspective, performing selection and configuration of workflows for responding to information and coordination requests while taking into account privacy implications for their execution. Figure 1 provides a graphical depiction of this concept.




Figure 1. PRIME concept: privacy-aware request processing

The privacy technologies used currently within PRIME include *searchable encryption*, *secure multi-party computation*, *function secret sharing*, and *differential privacy*. To complement those technologies, PRIME also leverages information-theoretic characterizations of what is being revealed through data releases, thus enabling informed decisions regarding the implications of sharing.

Much attention in the security community and the media is focused on adversaries that access data systems without authorization, for example by stealing credentials or exploiting software vulnerabilities. In this work, we focus on adversaries that access information solely through authorized channels. We note that such adversaries may still exploit that legitimate access to obtain information that an owner would prefer to keep private, for example by performing more data accesses than typically expected, or by inferring connections between data that are not explicitly related.

The remainder of the paper is organized as follows. Section 2 summarizes our HADR use case. Section 3 presents the core our PRIME information mediation platform, covering policy, workflow management, and services. Section 4 describes the privacy technologies that incorporated into PRIME to date and provides examples of their use. Section 5 describes directions for future work. Section 6 presents our conclusions.

2. Use Case: Humanitarian Aid/Disaster Relief (HADR)

Within our HADR scenario, a typhoon has caused extensive damage across a set of countries in the Pacific and relief (food, medicine, water, fuel, shelter, security, etc.) is needed in a number of communities. Adding to the complexity of the situation is the outbreak of a deadly and highly infectious disease that begins working its way through the populace.

The use case focuses on three fictional nations (Cebu, Bohol, Siquijor) that have sustained significant damage. To enable the use of state-of-the-art mapping and visualization capabilities, we elected to ground these fictional nations in real-world geographic entities, namely islands in the Philippines (see Figure 2). Each nation has five communities, which are marked by push-pins in the map.

HADR activities are being organized at multiple levels. Response Coordinators have been defined for each community and nation; there is also an over-arching International Response Coordinator to address cross-nation issues. There are ships from multiple nations in the general vicinity that could potentially provide resources to assist the impacted areas. An ad hoc coalition forms among nations (both aid providers and aid recipients) with varying degrees of amity/hostility/trust towards each other in order to distribute resources and to support necessary evacuations. Coalition members, while eager to coordinate on the relief effort, must take steps to ensure that information sharing is deliberate and conducted in a manner consistent with their organizational policies on information sharing.





Figure 2. HADR scenario setting: three fictitious island nations, each with five communities

To simulate the pandemic outbreak, we used the well-known Susceptible, Infectious, Removed, (SIR) compartmental model of disease progression [Kermack & McKendrick, 1927], augmented to support a Deceased compartment (i.e., an SIRD model);


- $S(t)$: # individuals not yet infected at time t
- $I(t)$: # individuals currently infected at time t
- $R(t)$: # individuals recovered at time t
- $D(t)$: # individuals deceased at time t

Possible state progressions for an individual are summarized below:



The standard model computes aggregate SIRD totals; we enhanced the model to track disease state for individuals and to support inter-community transmission. Figure 3 shows the equations used in the simulation along with a sample progression of the disease for a given set of model parameters.

Within this overall HADR use case, our focus to date has been on the three detailed scenario threads summarized below.



- $\frac{\partial S}{\partial t}(t) = -\beta S(t)I(t)$
- $\frac{\partial I}{\partial t}(t) = (\beta S(t) - k - d)I(t)$
- $\frac{\partial R}{\partial t}(t) = k(1 - d)I(t)$
- $\frac{\partial D}{\partial t}(t) = kdI(t)$
- $N = 5000$: population size
- $\beta * N = 0.5$: transmission rate
- $k = 0.2$: removal rate
- $d = 0.1$: mortality rate
- Initial infection rate: 10% in Maria, Siquijor

Figure 4. Model used to generate pandemic data, with graph of a sample progression for one community

- *Privacy-aware COP*: The operational objective in this thread is to provide a continuously updating common operational picture (COP) to the coalition members for ships in the area of responsibility. The privacy challenge is to provide information that will facilitate situational awareness and coordination without revealing information about ship positions, trajectories, and capabilities to parties that should not receive it. In particular, different nations will receive different views of the COP, based on controls imposed by the individual data owners.
- *Pandemic*: The operational objective in this thread is to predict the progression of a major disease outbreak through the impacted communities and to take steps to counter it. This thread introduces the challenge of protecting personally identifying information (PII) within medical records of individuals in the impacted communities while providing access sufficient to enable accurate characterization of the disease and its spread. A second privacy challenge relates to protecting information about certain aspects of the disease itself, to avoid inducing panic that could lead to mass migration and increased transmission among communities.
- *Aid Distribution*: This thread focuses on allocating and distributing resources (food, water, medicine) from coalition ships in the area to provide relief to hard-hit communities. Allocation and distribution planning require knowledge of ship positions, capabilities, and content, as well as of transportation and logistical capabilities in areas to which aid will be delivered. Each of these elements has contextually dependent privacy implications.

3. PRIME Platform

3.1 Overall Design

PRIME has been designed as an information mediator that provides access to data in accord with privacy restrictions imposed by data owners (see Figure 1). The PRIME platform leverages two proven, core technologies:

- A service-oriented architecture (SOA), called SIMON (Smart Integration Manager Ontologically Networked), that provides industry-standard identity management, policy enforcement, and micro-service capability integration. SIMON has been used to build and deploy a number of U.S. government systems for multi-nation information sharing.¹
- An adaptive agent platform, called Lumen, that is used as a high-level workflow engine for processing requests within the system. Lumen is a hardened implementation of the SPARK framework [Morley & Myers, 2004], which has been operationally deployed to support adaptive task execution within the U.S. Army's Command Post of the Future [Myers et al., 2011].

Requests are processed initially in SIMON, making use of its native authentication and logging capabilities in a pre-processing phase. Acceptable requests (as determined by identity management policies) are then forwarded to Lumen, which applies workflow models to respond appropriately. These responses can involve posting service requests back to SIMON for retrieving and processing data, or invoking various privacy technologies. Before making these service requests, the system consults a *policy reasoning engine*, built on an ontology framework called Sunflower [Ford et al., 2016], to determine appropriate controls on query-related data accesses. The Lumen workflow orchestration assembles the results and returns them to the user's display, via the SIMON framework services. Throughout, logging is performed to track all accesses and transformations to data, providing the means to support continuous awareness of what information has been released, to whom, and for what purpose.

The computational environment in which the system operates include a range of data and processing capabilities, wrapped as services within the SIMON SOA, that are leveraged during workflow execution. Most interesting here are the privacy-enhanced data and services, which are used to protect information. Section 4 elaborates further on those technologies.

¹ For example, SIMON was used to build the Cooperative Situational Information Integration (CSII) system for US Southern Command (US SOUTHCOM). CSII integrates partner nations with the U.S. into a regional, web-based, unclassified, network-centric information sharing system that spans air, maritime, and land domains.




Figure 5. PRIME platform architecture

3.2 Workflow Management

The workflow manager in the PRIME platform responds to user requests in a manner that is consistent with privacy policies defined by data owners. In particular, a given request may require a sequence of information retrieval and processing steps, each of which must be performed in a manner permitted by the policies. Thus, while the policy reasoning engine provides the capability for representing and reasoning with policies, the workflow manager is responsible for policy enforcement.

Workflows are used for two purposes in the current system. One is to support processing and coordination tasks that necessarily involve multiple data accesses and computations. For example, in response to a user's request for berth allocations, the PRIME workflow manager orchestrates predefined queries and computations to perform the allocation.

The second use of workflows is to compensate for technical shortfalls in query support within the encrypted database technology, which currently supports a restricted subset of SQL.² Relative to our HADR use cases, current gaps relate primarily to advanced/aggregate query capabilities. For such gaps, workflows decompose the complex queries into simpler ones that fit within the

² We anticipate less need for compensation of this type in the future, given the rapid advances being made in encrypted database technologies.

capabilities of the current query language, invoke those simpler queries, and then aggregate the results for presentation to the user. In other words, the workflow manager effectively performs the database operations not supported currently in the encrypted database. This approach has the potential to leak information from intermediate results. One possible way to address this leakage (to be explored in future work) is to compile this workflow into code that can be migrated to the client; if intermediate results are sent encrypted, then the potential for additional leakage is greatly reduced.

3.3 Policies

Data owners can define policies that limit access to structured data based on characteristics of the requester, request history, and request details. Policies are enforced through a query rewriting mechanism (described in the next section) that guarantees all information releases from encrypted databases are sanctioned by relevant data owners. Currently, policies are limited to controlling query access to structured data and web services. In future work, they will be extended to control access to a broader range of data types and information services.

The details of the policy language are beyond the scope of this paper. At its core, however, the representation allows expressive specification of constraints for accessing specific pieces of data as well as aggregate information (e.g., counts and averages), leveraging an underlying ontology of classes and relations. The policy representation also contains two constructs that are particularly important for privacy controls within coalition settings. One construct is an *override mechanism*, which enables one policy to take precedence over another in the event that they conflict. The second is an explicit linkage to organizational structures, which can provide the basis for defining overrides. Within the HADR use case, for example, nation-level policies are set to override community-level policies. Together, these constructs enable increased modularity of representation for policies within hierarchically structured organizations, compared to having to explicitly embed override conditions within policies for lower-level organizations. Overall, the policy representation is much richer than that of entitlement mechanisms in standard information systems, providing the flexibility to express the kinds of complex privacy restrictions necessitated in a coalition setting.

One interesting property of the policy reasoner is that it can generate *residual constraints* that serve as conditions for accessing requested information. In particular, policies do not gate access on a yes/no basis. For example, a policy may allow a Response Coordinator to access demographic information for people in a community but only for people older than thirteen. In the event that a request is made by the Response Coordinator for demographic information, he would be returned only the records for appropriately aged individuals.

Figure 5 illustrates the application of policies within the HADR use case, showing different views of information depending on the policies in force for different information requestors. Here, pie-charts depict percentages of the population in the different SIRD compartments. As shown, the International Response Coordinator is allowed to see nation-level views of the SIRD data (left); the Cebu City Coordinator is allowed to see a community-level view for its own community (middle); and the Bohol Nation Coordinator can see the community-level view for Bohol communities but only nation-level views for other communities (right).




Figure 6. Policy-differentiated information access, with different content and resolution based on community, national, and international roles

3.4 Policy enforcement

As noted above, the policy reasoner stores policies and makes decisions about what policies apply (and how) to a particular request. Actual enforcement of the policies, however, is done within the workflow manager.

Our current approach to policy enforcement centers on the notion of *policy-safe queries* (depicted in Figure 6). A given query is made policy safe by reformulating it to ensure that all information accesses are allowed given current policies. The simplest way to make a query policy safe is by reducing it to a ‘null’ query. However, our goal is to maximize the exchange of information while remaining compliant with policies. To this end, we developed an approach for automatically rewriting SQL queries to provide maximal access. With this approach, an initial query gets mapped to a collection of derived queries, some with additional WHERE clauses (to limit access to records) and some with SELECT clauses removed (to prevent access to data that should not be revealed to that user). The results of executing this modified set of queries are then merged to provide the overall policy safe response to the query. The additional constraints for the WHERE clauses and the set of SELECT clauses to be removed are generated by the policy reasoning engine.




Figure 7. Policy-safe query reformulation

Figure 7 shows an example of this policy safe reformulation for a particular SQL query that seeks to retrieve information (name, gender, birthdate, nation) for evacuees that have tested positive for an emerging virus, ZV1. Extant policies from different nations limit what information that requestor is allowed about their citizens: while Japan places no restrictions, the U.S. disallows access to birthdates and New Zealand disallows access to gender.

Consider first the case of Japanese evacuees. Because there are no restrictions in this case, the query can be made mostly as is, simply adding a constraint to the WHERE clause limiting the scope to evacuees from Japan (step 1). In step 2, separate queries are made for the countries that imposed limits on access. For records from New Zealand, the query is modified to exclude gender from the SELECT clause; for records from the U.S., the query is modified to exclude birthdate. In step 3, these partial results are then aggregated by the workflow manager into a composite result to be returned to the requestor.

Query: retrieve info for ZV1-positive evacuees

```

SELECT name, gender, birthdate, nation
FROM evacuee
WHERE ZV1Positive = TRUE;
  
```

Policies

- US: no access to birthdate
- NZ: no access to gender
- Japan: no restrictions

Policy-safe response:

1. Collect results for countries with no restrictions

```

SELECT name, gender, birthdate, nation FROM evacuee
JOIN community on community.id = evacuee.community_id
JOIN organization on organization.id = community.organization_id
WHERE ZV1Positive = TRUE AND organization_name IN ('Japan')
  
```

2. Collect filtered results for countries with restrictions

```

SELECT name, gender, birthdate, nation FROM evacuee
JOIN community on community.id = evacuee.community_id
JOIN organization on organization.id = community.organization_id
WHERE ZV1Positive = TRUE AND organization_name = 'NZ'
  
```

```

SELECT name, gender, birthdate, nation FROM evacuee
JOIN community on community.id = evacuee.community_id
JOIN organization on organization.id = community.organization_id
WHERE ZV1Positive = TRUE AND organization_name = 'US'
  
```

3. Aggregate results

Figure 8. Example of policy-safe query reformulation, with modifications highlighted in red.

The workflow manager also provides a *query relaxation* capability that can enable responses with lower fidelity than what was originally requested in response to policy restrictions, thus adhering to privacy requirements while releasing the maximally allowed amount of information. For example, if policies prohibit SIRD queries at the community level, the system will retry the query at the nation level. This type of relaxation is possible in cases where there is a natural generalization to a lower-fidelity characterization of the requested data.

3.5 Common Data Model

Following our initial platform development, we identified the need for a general data model that abstracts from particular data storage and processing representations to facilitate independent advances in design of the data processing, policy reasoning, and data storage specific schemas. A general model was needed that would both capture the ontologies of our scenarios and support automated mappings to the languages or schemas used for policy reasoning, request processing, and data storage. This model would also capture the meta-properties of the associated data and policies, including owner, time of creation, and their relationships. Such a model would greatly facilitate extension of the PRIME platform to other enterprise settings, which typically would have pre-existing data models.

To this end, we developed a *common data model* (CDM), using the Web Ontology Language (OWL) standard [Patel-Schneider et al., 2004], to represent relevant object classes, their properties, and relationships. The CDM is used to capture the information and relations needed to map data request representations to any independent data sources. It is also used as the underlying ontology for the policy representations. In this way, the CDM provides the semantic glue to connect the various components in the system, enabling metadata about information and services in the system to be captured in one place and then distributed to the modules that need it. It also isolates the policy engine and interface development from the specifics of individual data stores. Our first implementation handles SQL access to secure databases. Future developments will add mappings for RESTful data web services and unstructured file systems.

We developed automation tools to leverage the CDM representation for workflow and database accesses, as well as integration with the policy reasoner, enabling a common ontology to be used for data requests, processing, policy reasoning, and data accesses. This metadata explicitly captures the correlation intention of requested data (how requested data elements are to be connected in a single request) so that the appropriate set operations (e.g., JOINS in SQL) can be determined automatically from the mapping information associated with the CDM.

Figure 8 shows a subset of the HADR pandemic thread ontology and associated schema. The CDM-based classes and properties on the left are captured explicitly in OWL. The associated SQL schema, on the right, includes typical relational normalizations. The CDM intentionally abstracts the objects and instances within a particular domain, such as *person* or *nation*, as well as the hierarchy of their relationships such as *DiseaseStatus* as a specialization of *MedicalStatus*, which is in turn a specialization of *MedicalInformation*. While these abstractions are useful for building an ontology for policy representations, a typical normalized database schema is quite different in both structure and naming conventions. These differences were intentionally captured in our data models to explore the challenges in providing privacy-preserving technology for

existing enterprise systems. The needed mapping information is captured directly in the CDM model syntax and is then used to convey what information is being requested for determination of applicable privacy policies and request construction, as well as how to map this ontology to the SQL schema stored in the secure database for data access.




Figure 9. Pandemic subset of the HADR CDM with associated database schema

Figure 9 shows a sample mapping for the pandemic thread from the requested data elements represented in the OWL/RDF representation used by the CDM to the associated SQL query. The syntax of this CDM request contains the data elements requested using abbreviated RDF namespaces (e.g., prime#, medical#), as well as any user or policy-safe constraints that must be applied. In this example, the information requested is all persons’ first and last names, gender, resident community, and current disease state. The constraints include a user-defined filter for only those persons whose last checkup is before 10 April, 2017, and policy-based restriction that limits the request to data owned by the person’s community policy authority. The CDM-based mapping metadata is used by the automation tools within the PRIME platform to generate corresponding SQL queries.



Figure 10. Sample (abbreviated) CDM property chains and constraints auto-mapped to SQL

4. Application of Privacy Technologies

In this section, we discuss the various cryptographic and reasoning technologies used in PRIME to achieve privacy objectives for the HADR use case, focusing on their utility within the coalition setting and current limitations in their use.

4.1 Secure Multiparty Computation for Resource Allocation

Secure multi-party computation (MPC) enables agreed-upon computations to be performed on data supplied by multiple stake-holders without revealing any more information than the output of the function itself. In particular, the stake-holders provide encrypted versions of their inputs to the computation and derivative information from intermediate steps of the computation remain secret [Yao, 1982; Goldreich et al., 1987; Chaum et al., 1988; Ben-Or et al., 1988;]. One recent example of an application of MPC technology is for performing a secure probability analysis of satellite collision [Hemenway et al., 2016].

One use of MPC within PRIME is to support resource allocation. As an example, consider the following aid distribution task. There are k communities each of which requires M_c amount of a particular resource (e.g., food, medicine), and p resource providers each of which can supply R_i amount of the resource. The resource allocation task requires finding a set of feasible assignments

$$\begin{array}{ll}
 \text{Provider 1:} & M^1_1, \dots, M^1_k \\
 \text{Provider 2:} & M^2_1, \dots, M^2_k \\
 & \dots \\
 \text{Provider } j: & M^j_1, \dots, M^j_k
 \end{array}$$

such that

- for each provider $1 \leq i \leq p$: $M^i_1 + \dots + M^i_k \leq R_i$
- for each community $1 \leq c \leq k$: $M^1_c + \dots + M^p_c \geq M_c$

In particular, we require a given resource provider to completely service the request from an individual community; however, the provider can service multiple communities provided it has sufficient resources to satisfy their aggregate need. We also assume that provision of resources has an associated cost, which could be the cost of delivering the requested resources, or simply the motivation for servicing a particular request. Costs are considered private, with each participant seeking to protect this information from other coalition members.

We explored a range of different MPC algorithms to address this problem, considering options for optimal usage (relative to the provided cost models) and fair allocation that seeks to balance the costs evenly between the resource providers. Even within an MPC setting, care must be taken to avoid unintentional leakage of information. Consider an approach that seeks to optimize the allocation relative to the stated cost model by always selecting the lowest cost bid. In the case where there are two resource providers, each assignment of a resource provider to a requester reveals the relative costs for each of the two parties (i.e., the lower-cost bid always gets assigned). For this reason, we chose to mask this cost information by selecting the lower-cost bid only P percent of the time, for some selected threshold P . Through this probabilistic selection, the involved parties cannot be certain whether the assigned nation was in fact the lower bid for a particular request.

4.2 Searchable Encryption Meets Secure Multiparty Computation

A *searchable encryption* scheme securely encrypts data in a way that preserves one or more properties of interest, such as relative order or equality [Song et al., 2000; Boneh et al., 2004; Curtmola et al., 2006; Bösch et al., 2014]. PRIME provides a privacy-preserving relational database functionality implemented using the complementary technologies of secure multi-party computation (described above) and searchable encryption.

Unfortunately, while such encryption schemes may enable fast data access, they typically allow information leakage that may be observed by adversaries. In contrast, secure multi-party computation typically leaks very little information but is often several orders of magnitude slower than computation in the clear. In the Jana privacy-preserving database used in PRIME, relational queries written in SQL are answered in part by normal queries over such searchable encryptions, and in part by operations executed using a secure multi-party computation engine [Damgard et al., 2012]. By combining the two, and by allowing for each attribute in each database relation to be encrypted in one of several ways, Jana supports bespoke trade-offs between information leakage and query performance. Some configurations of Jana may be practically limited to 5,000 or so records in often-accessed relations. Other configurations may provide practical performance for much larger relations. Jana provides a tool for studying such trade-offs in practical use cases, something not previously reported in the secure computation literature.

4.3 Secret sharing

Secret sharing [Shamir, 1979] is a secure multiparty computation technique in which shares of some secret value are distributed to a group of participants by a trusted party. The original secret can be reconstructed only when a threshold number of shares are recombined. *Function secret sharing* (FSS) [Boyle et al., 2015] extends that concept to the computation of a function. Shares of a function are distributed to multiple participants in such a way as to enable them to each compute part of the overall function given a sufficient number of shares (otherwise nothing is revealed), and the result of the function on the secret inputs is additively recovered at the end of the computation without revealing anything about the inputs.

We use FSS to compute multi-level aggregation of SIRD population data, in accord with policies. In particular, the numerical SIRD pandemic population counts are stored over multiple function secret sharing services, along with representative attributes for functional processing. Our PRIME system, using the same workflow and policy decision mechanisms described above, accesses the policy-safe level of aggregations of these statistics (i.e., either nation-wide aggregation or community-wide aggregation) on a specific date. In this way, the FSS services, potentially kept by members of a non-trusting community, cannot reveal anything about the shared data without a minimum number of participants answering the policy-safe request.

4.4 Differential Privacy

Differential privacy [Dwork, 2006] enables statistical queries (in particular, aggregations) over a database of values, while minimizing the chances of identifying any individual within its records. Differential privacy is achieved through the use of principled noise injection for the results of aggregate queries to obscure the presence or absence of individuals within a database, with the amount of noise added linked to a specification of the degree of privacy that is to be maintained.

As a concrete example, consider an epidemiologist working to track and predict the disease spread. Policies are in place that enable him to access gender, birthdate, and disease status from the collected medical data as part of this process; however, to protect personal privacy, he is not allowed access to names. Consider if the epidemiologist is allowed access to fully accurate counts of people in the various SIRD categories. Suppose an update to the data is made to record a new entry for an infected individual. Queries immediately prior to and after the update would show that the new entry increased the Infected count by one, hence the addition must be infected.

Although the epidemiologist does not have direct access to the identify of this person, prior work has established that birthdate, gender, and zipcode are sufficient to identify individuals in 85% of cases [Sweeney, 2000]. As such, the epidemiologist could now likely determine the SIRD status as an indirect result. If instead the epidemiologist can access only differentially private counts of the number of infected people, the difference between the pre- and post-update count queries would not reveal the infected status of the new entry. Importantly, though, the differentially private counts can still provide useful statistical information for epidemiological modeling purposes [Ellis et al., 2017].

Differential privacy can also be combined with the multiparty computation technologies described above. As such, these enhancements have been included by the developers of the encrypted database and the FSS services in their technologies to further protect data privacy data.

4.5 Quantitative Leakage Analysis

To complement the above technologies for privacy preservation, we are exploring the use of technologies that can characterize what secondary information is being revealed through data releases. Awareness of this “leakage” is important for enabling informed decisions regarding the implications of sharing.

As noted above, all information releases and processing in the PRIME system are logged. This information feeds a real-time *quantitative information flow* analysis [Mardziel et al., 2013] that maintains an information-theoretic model of how accumulated releases decrease the receiver’s uncertainty regarding some target value that ideally would be kept private (such as a resource capacity). In particular, the model quantifies this leakage in terms of the number of bits of information that have been revealed.

To make this concrete, consider the task of allocating B berths in aggregate from nearby ships to transport seriously injured people from one of the impacted communities that has been particularly hard hit by the typhoon. The berths need to be available before some deadline T . Nations with ships in the area are willing to help out but do not want to reveal critical information about their ships (e.g., number of available/filled berths, position, travel speed).

The workflow required to complete the allocation can require multiple requests to the ship owners regarding their capabilities. One approach (which we call *separate*) first determines which ships can arrive by the deadline. For those ships, a form of binary search is then performed that establishes lower and upper bounds on each ship’s berth capacity. To start, each ship is asked whether it can provide the required B berths. If one answers positively, then it can be selected and the task is done. If not, the initial bounds are $[0, B]$ and the ships are then asked whether they can provide $B/2$ berths. If two of them can, the allocation is done. Otherwise, the bounds are adjusted: for a ship answering affirmatively, the lower bound is updated to $B/2$ while for ships answering negatively, the upper bound is reset to $B/2$. This process continues until a set of ships is determined whose sum of lower bounds exceeds the required number of berths. A variant algorithm (*combine*) merges the deadline and capacity queries. In this case, a negative response could mean either that the ship is too far away or that it lacks the requisite capacity, so can reduce the amount of information leakage.

These algorithmic variations (*separate*, *combined*) can lead to different allocations and leakage. Figure 10 illustrates this point for two separate berth allocation tasks (i.e., initiated for different communities), showing overall allocations along with leakage from the perspective of the USNS ships. In the case on top, no USNS ships are required and therefore drop out of the allocation process early with low leakage. In the case on the bottom, the capacity of the USNS is needed, incurring greater leakage as the binary search process refines the bounds on available capacity.




Figure 11. Comparison of leakage (in bits) for the combined and separate allocation algorithms

5. Future Challenges

Many issues remain in the area of technological support for privacy in coalition settings. Here, we describe three that we are addressing in our ongoing research.

- *Unstructured, streaming data*: As noted above, much of the information that is used in operational settings is unstructured (text, imagery, acoustic, PowerPoint, etc.) and continuously changing. In contrast, most of the work on privacy has focused on protecting information stored in structured databases. Streaming, unstructured data introduces significant challenges for developing mechanisms to minimize unnecessary sharing and to understanding the implications of information releases.
- *Privacy vs utility*: In this paper, we have focused primarily on mechanisms for safeguarding privacy. We are also exploring the tradeoff between increasing privacy versus decreased utility. For example, the work in [Ellis et al., 2017] examines the impact of different degrees of differential privacy on the accuracy of disease models that can be estimated by an epidemiologist.

- *Organizational policies:* Our work to date on policies has focused on controlling access to data. However, policies will also need to be formulated to control other aspects of the information systems. For example, policies should be defined to characterize who is authorized to write policies, or for what sorts of encryption levels are required to protect certain types of data.

6. Conclusions

This paper summarizes an approach to applying state-of-the-art privacy technologies to enable increased and informed information sharing within coalitions. Given the focus on extended collaborations within coalitions, we adopted a process-oriented perspective, performing selection and configuration of workflows for responding to information and coordination requests while taking into account privacy implications for their execution.

Appendix A: Privacy Technologies

The table below summarizes the technologies being used within PRIME and the organization that is providing them.

Organization	Technology
Galois	Searchable encryption Differential privacy Secure multiparty computation Quantitative leakage analysis
Pacific Science and Engineering Group	User workflow and policy modeling
SRI International	Policy representation and reasoning Workflow orchestration
Stealth Software Technologies	Function secret sharing Secure multiparty computation

Acknowledgments

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA), Air Force Research Laboratory (AFRL), and Space and Naval Warfare Systems Center, Pacific (SSC Pacific) under Contracts No. N66001-15-C-4065, N66001-15-C-4069, N66001-15-C-4070, N66001-15-C-4071, FA8750-16-C-0011, and FA8750-16-C-0022. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA, the Department of Defense, AFRL, or SSC Pacific.

The authors thank their colleagues who have contributed to the concepts and software described in this paper: Khang Duong, Thomas Garvey, Joshua Gould, Thomas Lee, and Laura Tam (SRI International) for their contributions to the PRIME platform; Linda Briesemeister, Samuel Wood, and Daniel Elenius (SRI International) for their contributions to the HADR CDM, policies, and the policy decision engine; Paul Bunn and Yuval Ishai (Stealth) for their work on MPC algorithms and their application to resource allocation; Jonathan Daugherty, Thomas DuBuisson, and Tristan Ravitch (Galois) for their contributions to the Jana privacy-preserving relational database technology; Jose Calderon (Galois) for his contributions to the quantitative leakage analysis capability; and Mark St. John (Pacific Science and Engineering) and Dusko Pavlovic for discussions related to user workflow and policy management.

References

- Archer, D. S., Bogdanov, D., Pinkas, B., and Pullonen, P. (2016). Maturity and performance of programmable secure computation. *IEEE Security and Privacy*, 14(5), 48–56.
- Ben-Or, M., Goldwasser, S., and Wigderson, A. (1988). Completeness theorems for non-cryptographic fault-tolerant distributed computing. *Proc. of the 20th Annual ACM Symposium on Theory of Computing*, 1–10.
- Bösch, C., Hartel, P., Jonker, W., and Peter, A. (2014). A survey of provably secure searchable encryption. *ACM Computing Surveys*, 47(2), Article 18.
- Boneh, D., Di Crescenzo, G., Ostrovsky, R., and Persiano, G. (2004). Public key encryption with keyword search. *Proc. of EUROCRYPT*, LNCS 3027, 506–522.
- Boyle E., Gilboa N., and Ishai Y. (2015). Function secret sharing. In: Oswald E., Fischlin M. (eds.) *Advances in Cryptology - EUROCRYPT 2015*. Lecture Notes in Computer Science, vol 9057. Springer, Berlin, Heidelberg.
- Chaum, D., Crepeau, C., Damgård, I. (1988). Multiparty unconditionally secure protocols. *Proc. of the 20th Annual ACM Symposium on Theory of Computing*, 11–19.
- Curtmola, R., Garay, J., Kamara, S., and Ostrovsky, R. (2006). Searchable symmetric encryption: improved definitions and efficient constructions. *Proc. ACM Conference on Computer and Communications Security*, 79–88.
- Damgard, I., Keller, M., Larraia, E., Pastro, V., Scholl, P., and Smart, N. (2012). Practical covertly secure MPC for dishonest majority – or: breaking the SPDZ limits. Cryptology ePrint Archive: Report 2012/642.
- Dwork, C. (2006). Differential Privacy. *Proc. of the 33rd Intl. Colloquium on Automata, Languages, and Programming*, Vol. 4052. Springer, 1–12.
- Ellis, T., Lee, T., Lepoint, T. and Myers, K. (2017). On the usability of differentially private SIRD data during a pandemic outbreak. *Proc. of the 13th Symposium on Usable Privacy and Security* (Poster Track).
- Ford, R., Denker, G., Elenius, D., Moore, W., and Abi-Lahoud, E. (2016) Automating financial regulatory compliance using ontology + rules and Sunflower. *Proc. of the 12th Intl. Conference on Semantic Systems*, Leipzig, Germany, 113–120.
- Goldreich, O., Micali, S., and Wigderson, A. (1987). How to play ANY mental game. *Proc. of the 19th Annual ACM Symposium on Theory of Computing*, 218–229.

- Hemenway, B., Lu, S., Ostrovsky, R., and Welser IV, W. (2016). High-precision secure computation of satellite collision probabilities. *Proc. of the Intl. Conf. on Security and Cryptography for Networks*, 169–187.
- Kermack, W.O. and McKendrick, A.G. (1927). A contribution to the mathematical theory of epidemics. *Proc. of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 115 (772): 700.
- Mardziel, P., Magill, S., Hicks, M. and Srivatsa, M. (2013). Dynamic Enforcement of Knowledge-based Security Policies using Probabilistic Abstract Interpretation. *Journal of Computer Security*, 21:463–532.
- Morley, D. and Myers, K. (2004). The SPARK agent framework. *Proc. of the 3rd Intl. Joint Conf. Autonomous Agents Multi Agent Systems*.
- Myers, K., Kolojechick, J., Angiolillo, C., Cummings, T., Garvey, T., Gervasio, M., Haines, W., Jones, C., Knittel, J., Morley, D., Ommert, W., and Potter, S. (2011). Learning by demonstration technology for military planning and decision making: a deployment story. *Proc. of the 23rd Conf. on Innovative Applications of AI*.
- Patel-Schneider, P., Hayes, P. and Horrocks, I. (2004). OWL web ontology language semantics and abstract syntax. W3C Recommendation.
- Shamir, A. (1979). How to share a secret. *Communications of the ACM*, 22 (11): 612–613.
- Song, D., Wagner, D.A., and Perrig, A. (2000). Practical techniques for searches on encrypted data. *IEEE Symposium on Security and Privacy*, 44–55.
- Sweeney, L. (2000). Simple demographics often identify people uniquely. Carnegie Mellon University Working Paper. <http://dataprivacylab.org/projects/identifiability/paper1.pdf>
- Yao, A.C. (1982). Protocols for secure computations. *Proc. of the 23rd Annual Symposium on Foundations of Computer Science*, 160–164.

Knowledge Systems for Coalition Operations 2017

**The Role of Transactive Memory (TM) in Proactive Decision Support (PDS)
(Position Paper)**

Erica P. Viklund, Heather M. Oonk, and Manuela Jaramillo
Pacific Science & Engineering Group
9180 Brown Deer Road, San Diego, CA

Abstract

Advances in technology have exponentially increased the information and data at our fingertips. While there are many benefits of such access, a tradeoff is that information seekers can be overwhelmed by the vast sea of information at their disposal. Challenges multiply when information seekers operate as part of a team where there are differences in knowledge, information access, and decision-making responsibilities. Coalition operations are examples of such situations, involving decisions that impact a complicated network of different countries and actors. Proactive decision support (PDS) tools have the potential to make more manageable the tasks of selecting, verifying, compiling, and analyzing relevant information, so that good decisions can be made more efficiently. Effective PDS requires a system that “understands” and adapts to the context in which information seeking and decision-making occur. Context includes aspects of the physical environment within which the technology and user are embedded, and the cognitive or mission objectives of users. We argue that for teams, PDS context must also include a collection of team member and team dynamic variables such as shared and differential tasks, requirements, knowledge, and expertise. Collectively, these variables can be conceptualized as transactive memory (TM). We describe how PDS that incorporates TM variables as a form of context can facilitate and streamline validation and communication of information among team members, which is crucial for realizing the potential benefits of PDS for coalition operations. We discuss considerations for implementing TM variables into PDS tools and key research and development questions to be addressed.

A Case for Proactive Decision Support

Dateline: July 2017

A large, distributed, multinational team works against the clock to provide aid to a disaster-stricken region. Despite their best intentions and efforts, priceless time is lost in the team's attempts to coordinate with each other, to share relevant information, and to sift through irrelevant information and numerous unintentional "red herrings." As a result, they fail to make decisions in a timely and efficient manner, countless lives are lost, and much suffering ensues. The world resolves to do better next time...

Dateline: July 2027

A massive earthquake, volcanic eruption, and resultant tsunami cause mass destruction in Southeast Asia. Teams from across the globe mobilize to provide humanitarian assistance and disaster relief to the stricken local population. Despite the unprecedented scope and complexity of the operation, response teams can provide effective care and relief by using new proactive decision support (PDS) tools developed by the US Navy's Office of Naval Research (ONR). Sophisticated artificial intelligence-based PDS tools automatically recognize key decision events, gather decision-relevant information, and distribute it to those personnel responsible for making decisions. When specific expertise is needed, or new information becomes available, the PDS system automatically routes the information or requests to the correct people. When decisions are made, the system accurately predicts follow-on events and decisions, and prepares for them as well. Using the new PDS tools, the world relief organizations can take control of the disaster in an organized fashion, and effectively and efficiently provide aid when and where it is needed the most. Despite the massive scale of the disaster, the casualties are kept to a minimum, and the survivors can quickly rebuild their lives and prosper. In after action discussions about the response, the teams agree that their PDS tools were invaluable for the rapid, effective communication of information and sound decision-making that directly contributed to the success of their missions. Moving forward, smart PDS systems will be a core component of coalition operations...

How to Make PDS a Reality: Introduction

Advances in technology have exponentially increased the amount of information and data at our fingertips. While there are many benefits of such easy access, a tradeoff is that information seekers can be easily overwhelmed by the vast sea of information at their disposal. It can be difficult and time-consuming to sift through available information and to select, verify, compile, and analyze the information that is relevant. The consequences of these tradeoffs can be especially significant for those who seek information to make rapid, high-stakes decisions in complex environments. Further, the challenges multiply when information seekers are operating as part of a team when there are differences in knowledge, access to information, and responsibilities in the decision-making process. Coalition operations such as the humanitarian assistance and disaster relief efforts described above, are examples of such situations, involving decisions that may impact a complicated network of different countries and actors.

Proactive decision support (PDS) tools can improve the efficiency of the decision-making process, for example by automatically monitoring information feeds, recognizing and alerting users to decision events, and gathering and presenting critical information. Effective PDS requires a system that "understands" and adapts to the context in which information seeking and decision making occur. This decision context includes various aspects of the physical environment within which the technology and user are embedded [1] and the cognitive or mission objectives of users [2]. Overall (environmental or mission) context may also interact with information provided by various sources, including decision cues and other types of context, to influence the decision-making process, as shown in the notional *decision support structure* in Figure 1. The effect of context on the outcome of the decision(s) may be determined by the *decision factors* (type, timeliness, priority, consequences, etc.). Further, decisions made may influence other decisions and/or impact context.




Figure 1. Decision support structure with example details.

In this position paper, we argue that for teams, PDS must also be responsive to context relevant to team member and team dynamic variables such as shared and differential tasks, requirements, knowledge, and expertise [3]. Collectively, these variables can be conceptualized as a form of transactive memory (TM) – the concept of collective knowledge possessed by individual group members, and the shared awareness within the group of who knows what. We maintain that considering TM as a form of context and providing a mechanism to foster TM in PDS tools will facilitate and streamline communication of information among team members. We suggest that addressing TM variables can increase the potential performance and workload benefits of a PDS tool for complex, time-sensitive, coordinated team missions. The paper will discuss the following:

- The concept of TM and its relevance to coalition operations
- Considerations for a TM-supportive PDS tool
- Research and development challenges for TM in PDS

The paper concludes with a high-level summary of the potential value of a TM-supportive PDS system to coalition operations.

Transactive Memory

As noted, an effective PDS system must be *context-aware*, incorporating contextual elements such as the environment and mission, a user's available information sources, and factors specific to the decision at hand. For teams, an often-overlooked aspect of context is TM. Moreland defines TM as the sum of "knowledge possessed by individual group members with a shared awareness of who knows what" ([4], p.5). With accurate shared awareness, team members can better evaluate how and to whom to distribute new information, and from whom to request information. Shared awareness and differentiated knowledge (unique knowledge and expertise contributed by individual team members) form the *structure* of TM; there are also transactive *processes* by which team members cooperatively store, retrieve, and communicate information (cf., [5]). A *transactive memory system* (TMS) is characterized by the dynamic interaction of team members' individual TM to draw on and integrate knowledge to achieve a team task [6]. Performance benefits of a strong TMS include reduction in errors and better recall of task procedures (see [7]). TM research originally focused on two-person teams, but it has been extended to

larger groups and teams¹ with similar results [6]. PDS tools need the ability to represent, evaluate, and “perceive” TM context and supporting decision processes by adaptively facilitating effective TMS development and function.

TM in Coalition Operations. Coalition operations present a special type of team environment, one in which a robust TMS has the potential to provide immense benefits. Specific characteristics and requirements of coalition operations that research suggests are supported by TM include:

- The need to understand expertise and differentiated knowledge across distributed, cross-culture teams [4]
- The need for trust across unfamiliar team members (note that TM has been found to be both an antecedent to and an outcome of trust [8])
- Substantial involvement of transactive processes, such as collaboration, communication, and feedback [9]
- The need to understand team member culture and leadership styles [10]

At the same time, coalition operational environments are intrinsically suboptimal for the *development* of TM. Coalition teams are frequently ad hoc, and rarely involve the amount of group training required to form strong TMSs. In the humanitarian assistance and disaster relief scenario described above for example, teams from around the world who may never have worked together must rapidly coordinate and execute a set of missions with a complexity and scope that can never be fully trained. Further, with an effort this massive the amount of information and number of information sources would be completely overwhelming. Yet to make sound decisions, the credibility and relevance of every piece of information must be established, and the right information must get to the right people at the right time. In other words, a strong TMS is critical for efficient, effective decision-making. The coexistence of the need for and likely absence of TM provides a unique opportunity to leverage a TM-supportive PDS system.

Considerations for a TM-supportive PDS Tool

There is evidence that technologies such as information directories and repositories can help to build TMS within teams (e.g., [11]), but that these technologies fall short of producing expected performance benefits. Lewis & Herndon [9] argue that such technologies do not allow full development of a successful TMS because while they provide TM *structure*, they fail to support the transactive *processes* that are crucial to a TMS. A PDS tool offers additional capabilities beyond information directories and repositories in that it can be designed to model context, including transactive processes. This section discusses considerations for a PDS tool designed to facilitate strong TMS and improve performance of team tasks that require identification, selection, verification, compilation, analysis, and communication of information.

Transactive Memory Information Requirements for PDS. The specific information requirements of a TM-supportive PDS tool will be partially dependent upon the application for which it is being developed. However, there are some overarching considerations that will be relevant to the TM aspects of decision support in most any application. Because TMS is being conceptualized as a set of context variables that shape the decision-making process, there will be representations of TMS that are not communicated explicitly to users, but rather work “under the hood” to support the user and to help the PDS tool evolve and improve over time. However, there will be occasions in which there is explicit communication of TM information between the system and the user. For example, a new team member not familiar with the expertise of other team members may ask “Who would know local resources in small Southeast Asian villages that might have access to baby food for the families displaced by the tsunami?” Explicit communication between users and with the PDS tool will naturally be more frequent when the tool is in its initial stages and the underlying decision structures haven’t had an opportunity to “learn”. There will also be more instances of explicit communication of TM information when there are one or more new team members, as the baseline states for the new member(s) become instantiated in the PDS tool. Broadly speaking, the explicit communications will either *push* information to a team member, or *pull* information from a team member. When explicit communication is required, the design of the PDS tool interface should ensure that the communication happens as unobtrusively as possible. The goal is to allow the PDS tool to develop a TMS (an internal representation of who knows what) so that the tool can get the right information to the right people at the right time with minimal burden on the user. The information to be captured, understood, and represented in a TM-supportive PDS tool can be categorized in terms of the “five Ws”, as summarized in Table 1.

¹ A distinction can be made between ‘groups’ and ‘teams’, with teams typically having more clearly specified roles, tasks, and relationships, such as in military command and control domains. The terms are used together here because although the TM literature frequently addresses groups, many aspects of the work apply readily to teams (see [15]).

Table 1. Summary of general information requirements for development of a TM-supportive PDS tool.

TM Information	
Who	Which team member(s), if any, possess a specific piece of knowledge? Which team member(s), if any, need that knowledge?
What	What information category/categories are relevant to the current decision process? What specific piece(s) of knowledge are relevant to the current decision process?
Where	Where does a specific piece of knowledge reside? <ul style="list-style-type: none"> • In the expertise of a team member? • In an external repository such as a directory or database? • As tacit knowledge within the TMS?
When	At what time in the decision-making process should a specific piece of knowledge be communicated to a team member for <i>proactive</i> decision support?
Why	What is the relevance of a specific piece of knowledge to the current decision process?

Task Selection. There are certain types of tasks that stand to benefit the most from a strong TMS, and these tasks should be given special weight in the development of TM-supportive PDS tools. Lewis & Herndon [9] provide a useful discussion of relationships between TMS and task types. They characterize tasks based on processes and structural qualities. Task processes include “produce” (generate ideas), “choose” (select a solution to a problem), and “execute” (perform a set of actions to achieve a goal). Task structural qualities relate to whether the associated responsibilities are unitary or divisible among team members, whether the task goals are cooperative or competing across team members, and the degree to which there is a single, demonstrably correct task outcome. To summarize, they conclude that while TMS is important for these task types, the ones that stand to benefit most from a strong TMS are “execute” tasks with divisible responsibilities, cooperative goals, and a demonstrably correct solution. More so than other task types, efficient and effective team performance on tasks with these characteristics relies upon (a) a solid TM structure with diverse and specialized knowledge that is readily accessible from known, credible sources, and (b) well-developed transactive processes that facilitate communication, coordination, and integration of this knowledge.

For example, in the tsunami scenario above, an operationally-appropriate *execute* task may be “Provide medical aid to local villages that require it.” While each coalition member may be able to perform this task to some degree on their own, coordinating to divide the task to take advantage of each nation’s strengths/expertise would provide a more optimal solution. One nation may be able to provide the best medical staff and supplies, another may have the closest transport resources available, another may know the location of the most appropriate linguist to help the medical staff, and yet another may have knowledge from previous events related to the most efficient and safe route to use to get to local villages. A TMS-enabled PDS system would provide a mechanism to support this coordination and, thus efficient execution, of the task.

Representing TM Variables in a PDS Tool. Palazzolo et al. [12] discuss the utility of considering TM from a network perspective. Palazzolo et al. [12] developed and tested a network model of TM that illustrates how TMS could be instantiated as a form of context in a PDS tool. The model was based on a conceptual framework that specifies three interrelated transactive processes by which a TMS develops. These processes are directory updating (each team member’s dynamic understanding of who knows what within the team), communication to allocate information, and communication to retrieve information. The team’s initial attributes and the success of these transactive processes result in two measures of TMS development: accuracy in expertise recognition and differentiation of knowledge [12]. The starting attributes of the team that are captured in their model are initial knowledge, initial accuracy of expertise recognition, and network (team) size.

This network perspective allows not just individuals, but also the connections between them to be modeled and measured. As Stanton et al. [13] note in their related discussion of distributed situation awareness, the focus in a network is on *links* (interactions and transactions between team members) rather than on *nodes* (the information processing of individual team members). In the tsunami scenario, these links may be transactions between members as they share information to coordinate each step of the task (move resources using assets, coordinate with translators, plan safe land route, ...).

Research and Development Challenges for TM in PDS

Two important research and development considerations for a TM-supportive PDS tool are (a) validating the representation (capture, learning, and updating) of TM variables within the other decision structures of the tool, and (b) testing to determine whether the inclusion of TMS context does in fact add value to the PDS tool. The validation and testing process both require

measures of TMS. Traditionally, measurement of TMS has consisted primarily of inferences generated from team members' recall of information in task performance, observation and evaluation of team members' behaviors and communications during task performance, and team members' post hoc self-reports about the credibility of other team members' knowledge [6]. While these measures may be suitable for experimental settings, there are multiple constraints that limit their utility for measurement of TMS in field settings (see [6] for a discussion).

In an effort to overcome the limitations of the traditional TMS measures, Lewis [6] developed and validated a TMS scale that would be appropriate for field use. The results of the testing reported by Lewis [6] indicated that his fifteen-item scale is a valid measure of TMS that can be applied across tasks and teams. Another candidate measure could be derived from the Event Analysis of Systemic Teamwork (EAST) methodology described by Stanton et al. [13]. While designed for analysis of distributed situation awareness, the methodology should be adaptable for TMS constructs as well. While there is further testing to be done, Lewis's [6] scale and the EAST methodology [13] provide candidate field-appropriate TMS metrics that could be used as part of a suite of measures to evaluate TMS networks for PDS.

Additional measures could include automated data collection built into PDS tool prototypes, and surveys administered during coalition training exercises. In simulations using a TM-supportive PDS tool where the information available to a network can be controlled, the pieces of information that are propagated through the TMS network to and from each agent can be extracted and evaluated. Thus, additional candidate measures of TMS can be gleaned from using a signal detection framework [14] to conceptualize the TMS-dependent information collection, validation, and communication aspects of the decision-making process. Within the signal detection framework, TM performance could be captured by analyzing the relative frequency of "hits" (valid, relevant information is communicated to the right people), "misses" (valid, relevant information is NOT communicated), "correct rejections" (invalid, irrelevant information is NOT being propagated), and "false alarms" (invalid, irrelevant information is propagated). Once a TM-supportive PDS tool has had an opportunity to learn, it should maximize hits and correct rejections, and minimize misses and false alarms. If the tool was not learning as expected, the signal detection analysis would provide diagnostic information by allowing the identification and localization of the issue(s) (e.g., invalid information propagating between specific nodes).

Taken together, the resultant data would inform the design of next-generation PDS tools and the associated human computer interfaces. Further, the outputs of several of these measures can be used to help diagnose the source of shortcomings in under-performing teams, which may have implications for improvements to team composition and system design. For example, the relative scores across the three constructs that comprise Lewis's [6] scale can reveal weaknesses in team members' knowledge, their coordination of knowledge, and/or their perceptions of each other's expertise. Depending on the score profile, the weaknesses might be mitigated by a change in team composition or a change in training practices to improve TMS development.

While the test and validation process required for a TM-supportive PDS tool will be extensive, the extant body of work has established a strong conceptual framework and promising network architectures that can readily be integrated with other PDS decision structures [12]. Further, there is a set of candidate measures available that can provide converging data to help refine and validate TMS components of a PDS tool. There is a clear need for PDS in team tasks that require complex decision making, such as those commonly encountered in coalition operations. Because the differentiated knowledge and transactive processes among team members are such crucial drivers of decision making performance, TMS must be considered as an essential element of a PDS tool that will be successful in improving performance and reducing workload.

Discussion and Conclusions

At a time when vast amounts of information are available from myriad sources, PDS tools have the potential to make the task of selecting, verifying, compiling, and analyzing information more manageable so that decisions can be made more efficiently and effectively. A PDS tool that incorporates TM as a form of context can increase the value of a PDS tool for team tasks by also representing the distribution of knowledge across team members, team members' meta-knowledge about where information resides, and information about the credibility of information sources.

The benefits of TM to coalition operations can be illustrated by returning to the decision support structure discussed above in Figure 1. The fact that the mission is being conducted by a coalition team might be considered a type of overall context, and the team variables that stem from this, such as understanding/uncertainty about the distribution of knowledge, the need for collaboration, etc. might be representative of decision context. Figure 2 adds to the decision support structure several TM variables (shown in yellow and with bolded text) that can inform and improve decision-making in a coalition team context. The variables relate to the specific requirements of coalition teams discussed above that include support for 1) differentiated

knowledge understanding in distributed teams [4], 2) trust in unfamiliar team members [8]), 3) collaboration, communication, and feedback [9] and 4) team member culture and leadership styles [10]. Also shown in Figure 2 is the update and increase in TM as an outcome (impact) of the decision-making process.




Figure 2. Decision support structure for coalition operations, with TM variables (shown in red).

If properly designed and implemented, a PDS system that supports TM should improve operations by facilitating the sharing of relevant information among team members, thus reducing workload, increasing the efficiency of information gathering, verification, compilation, and analysis, and, ultimately, supporting optimal effective decision making. Thus, consideration of TM variables is crucial for the development of a PDS tool that will maximally support the complex, high-consequence decisions and large, disparate teams that characterize so many coalition operations.

References

- [1] A. K. Dey, G. D. Abowd and D. Salber, "A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications," *Human-Computer Interaction*, vol. 16, pp. 97-166, 2001.
- [2] J. Hong, E. Suh and S. Kim, "Context-aware systems: A literature review and classification," *Expert Systems with Applications*, vol. 36, pp. 8509-8522, 2009.
- [3] Pester-Dewan, M. R. A. J. and H. M. Oonk, "Knowledge flow in command & control: Developing intelligent aided communication. (Technical Report)," Pacific Science & Engineering Group, San Diego, CA, 2005.
- [4] R. L. Moreland, "Transactive memory: Learning who knows what in work groups and organizations," in *Shared Cognition in Organizations: The Management of Knowledge.*, L. L. Thompson, J. M. Levine and D. M. Messick, Eds., Mahwah, NJ, Lawrence Erlbaum Associates, 1999, pp. 3-32.
- [5] D. M. Wegner, T. Giuliano and P. Hertel, "Compatible and Incompatible Relationships," in *Cognitive interdependence in close relationships.*, W. Ickes, Ed., New York, Springer-Verlag, 1985, p. 253-276.
- [6] K. Lewis, "Measuring transactive memory systems in the field: scale development and validation," *Journal of Applied Psychology*, vol. 88, pp. 587-604, 2003.
- [7] R. Moreland and L. Myaskovsky, "Exploring the performance benefits of group training: transactive memory or improved communication?," *Organizational Behavior and Human Decision Processes*, vol. 82, pp. 117-133, 2000.
- [8] M. Ashleigh and J. Prichard, "An integrative model of the role of trust in transactive memory development," *Group & Organization Management*, vol. 37, pp. 5-35, 2012.
- [9] K. Lewis and B. Herndon, " Transactive memory systems: Current issues and future research directions," *Organization Science*, vol. 22, pp. 1254-1265, 2011.
- [10] A. M. Phebus, B. Gitlin, M. L. Shuffler and J. L. Wildman, "Leading Global Virtual Teams: The Supporting Role of Trust and Team Cognition," in *Cross-Cultural Interaction: Concepts, Methodologies, Tools, and Applications*, E. Khosrow-Pour, Ed., Information Resources Management Association, 2014, pp. 362-384.
- [11] S. Y. Choi, H. Lee and Y. Yoo, "The impact of information technology and transactive memory systems on knowledge sharing, application, and team performance: A field study," *Management Information Systems Quarterly*, vol. 34, p. 855-870, 2010.
- [12] E. Palazzolo, D. Serb, Y. S. C. She and N. S. Contractor, " Coevolution of communication and knowledge networks in transactive memory systems: using computational models for theoretical development," *Communication Theory*, vol. 16, pp. 223-250, 2006.
- [13] N. Stanton, P. Salmon and G. Walker, "Let the reader decide: a paradigm shift for situation awareness in sociotechnical systems," *Journal of Cognitive Engineering and Decision Making*, vol. 9, pp. 44-50, 2015.
- [14] J. A. Swets, "Detection theory and psychophysics: A review," *Psychometrika*, vol. 26, pp. 49-63, 1961.
- [15] S. Mohammed and B. Dumville, "Team mental models in a team knowledge framework: expanding theory and measurement across disciplinary boundaries," *Journal of Organizational Behavior*, vol. 22, pp. 89-106, 2001.

Cover Sheet

Symposium: 9th Knowledge Systems for Coalition Operations Conference

Track 10: Knowledge Systems for Coalition Operations

Authors:

Adrienne J. Raglin
U.S. Army Research Lab, RDRL-CII-B
2800 Powder Mill Rd, Adelphi, MD 20783
adrienne.j.raglin.civ@mail.mil

Christian D. Schlesiger
U.S. Army Research Lab, RDRL-CII-B
2800 Powder Mill Rd, Adelphi, MD 20783
christian.d.schlesiger.civ@mail.mil

Data Analyzer Software: a Knowledge System Supporting Coalition Partners Information Sharing

Abstract

As current and future operations integrate soldiers from multiple nations, information that supports short term and long term teaming is critical. Among coalition forces, it is important to maintain unity of effort, to plan concurrently, and to make adjustment in sync, ensuring operations are carried out successfully. Combatant commanders have many responsibilities including ensuring the capability and capacity of the forces with partnering nations. However, in multinational operations there is the added need to consider differences in organizations, doctrines, terminologies, and objectives. This can be achieved through knowledge capturing, information sharing, and training. Additionally, giving commanders required information with explanation, linking knowledge and uncertainty could improve teamed operations in complex and dynamic environments. The Data Analyzer is an adaption of training software previously developed at the Army Research Laboratory. Now that this training software is being used by coalition partners and the US the Data Analyzer has been expanded as a platform for knowledge capturing and information sharing. The Data Analyzer provides commanders with the ability to view detailed experiential knowledge and find trends in tactics, techniques and procedures (TTPs) employed within coalition partners. This information from the analyzer provides the coalition forces with highlighted similarities and differences that aid in coalition engagement preparation and insights into actions that can impact coalition mission TTPs. We present the Data Analyzer software and illustrations utilizing this approach in supporting knowledge capturing and information sharing for multinational operations.

1 - Introduction

“Multinational operations are operations conducted by forces of two or more nations...”[1]. These coalitions are formed for humanitarian aid and disaster relief, as well as, military actions against perceived threats. For these types of partnerships to be successful “unity of effort”, knowledge between partners and clear mission goals are some of the required elements. Interoperability between nations and standardization increase effective coalition operations. For coalition operations, information must be coordinated to fully understand tactics, integrate and synchronize missions, transfer and disseminate intelligence. Since these types of operations can have multiple chain of commands, decision making considerations must take into account the overall mission, the operational environment, size of the force, risk, duration and rules of engagement. In coalition engagements each country’s regulations mandate when, where, and

how information is shared. For military coalitions, “effective information sharing enables the DoD to achieve dynamic situational awareness and enhance decision making to promote unity of effort across the Department and with external partners [2].”

Taking a more military view, information sharing is “the willingness to provide others with information [3]” particularly incorporating information provided from experienced soldiers or experts. Commanders use information sharing to “enhance their ability to execute battle commands [2].” Using common technology coalition partners may be able to increase their effectiveness in sharing information. In this case, the goal of any technology is to “provide a common understanding of a shared vision, mission and governing principle.” The Applied Anomaly Detection Tool (AADT) and AADT Data Analyzer (DA) is a software system that aides in achieving this goal.

In the next section of this paper, we will present the AADT and AADT DA as an information capturing and information sharing software system that can be used by military coalitions and teams for training, situational awareness, and mission planning/reviewing as well as decision making. In section 4 we describe using AADT and AADT DA for information sharing among coalition commanders. In section 5 we describe using AADT and AADT DA for information sharing among partners’ units. In section 6 we describe how AADT and AADT addresses information sharing challenges. In section 7 we briefly describe future work.

3 – AADT and AADT DA

The AADT Data Analyzer is an application that was originally developed to organize and present the data collected by the AADT. AADT is a software program designed for testing and training soldiers in finding targets and areas of interests in military relevant operations. In previous studies conducted at Army Research Laboratory, targets and areas of interest were referred to in general as anomalies. Soldiers often described key indicators of potential threats as anomalies and this terminology was adopted during the development of the AADT software. The AADT software has two subsystems, one for the trainer and one for the trainee. The trainer subsystem is used to create courses and classes. It is also used to configure the ordering of subject matter expert (SME) feedback, to set access control of courses created by the trainer, and to manage trainee evaluations. The trainer subsystem gives the trainer the ability to generate and modify the content of material that can be used to create a course. The trainee subsystem is the component of the AADT that allows a class to be taken. The trainee subsystem shows background information, interactive training or testing modules, SME feedback, quizzes, and evaluations. Since the primary use of AADT is for visual search and detection, a series of images are displayed by the software of different targets at various distances, in various terrains and environments. With the ability to present SME feedback, expert knowledge and instructional guidance is provided. The training AADT enhances the situational awareness or the understanding of the operation battlespace.

4 - Using AADT and AADT DA for information sharing among coalition commanders

In the AADT software commanders from the different coalition partners, as well as across teams, can use the trainer subsystem to generate a series of modules to share information related to key concepts for missions or tasks, particularly for those involving visual search of potential adversarial environments. There are many features for creating, modifying, and organizing basic background information and scenario driven modules. The primary feature is the anomaly page option. Anomaly pages are used to upload any image (single view, multi-view, or panoramic) from any type of sensor. Individual or groups of commanders can then draw polygons on the image where anomalies are located. To standardize the type of anomaly (target and AOI), the type is selected from a menu in the interface of pre-defined types. After selecting the type, a dialog box allows metadata to be added further describing the anomaly shown in the polygon. These descriptions are called key feedback. These are unique for each polygon type selected. As previously mentioned, AADT provides a robust set of pre-defined types and properties that can describe a wide variety of anomalies. However, new ones can be added by selecting 'Add New' option from the drop-down lists and entering the additional information. The key feedback information can be copied to an image feedback area. The Image Feedback area allows commanders to provide additional information about the image. This additional information can be a discussion of the overall scenario represented by the image or by categories of items. The highlighted items can be grouped by polygon types and the additional information can be organized to match these categories. The purpose of this feedback is to annotate why the individual anomalies are important to understanding the situations the image represents. The images can be grouped in modules by a designated theme, site, target type, or mission route.

The AADT DA interface has several features that support information sharing at the commander level. In the 'View' option a single image or a collection of images with anomalies highlighted and feedback for review (Figure 1). Figure 1 shows image feedback which is a general description of the scene and key feedback which gives general descriptions of the AOIs highlighted with black polygons. Having this capability allows the methodologies used by the different forces for identifying potential threats and assessing the critical components of the environment that may call for heightened situational awareness.




Figure 1. 'View' option of AADT DA

In addition, AADT allows commanders' and soldiers' demographic information to be recorded and analyzed in AADT DA as 'Demographics' (Figure 2). Multiple entries can be viewed and accessed. This information aide in understanding the background of the partners, and how their information is shaped by the doctrine and practices set by their country's policies. Sharing this information can indicate strengths and challenges within the specific tasks by the partners. This can facilitate modification in assignments or actions that can improve unit level dynamics, increasing mission success.

Demographics		
1	ID Number: 7001	Age: 26-35
	Number of Years in Service: 0-5	MOS Primary Code: 68W
	Military Rank/Grade: PVT	MOS Primary Job Title: 68W10
	Land Navigation Qualifications?: Yes	MOS Secondary Code: None
	Weapons Qualifications: Marksman	MOS Secondary Job Title: None
	List Additional Weapons: NA	MOS Time (in years): 0-5
	Fluent Languages (exclude English): NA	Have you ever been deployed?: No

Figure 2. Demographic information

5 - Using AADT and AADT DA for information sharing among partners' units

After commanders use AADT and AADT DA to create and review modules related to various aspects of the operation then sharing that information at the unit levels can be done in several ways supporting training and missions. The motivation for commanders using AADT DA for information sharing include 1) enhanced dissemination of tactics, techniques and procedures (TTPs) that have been agreed upon for an event, 2) enhanced situational understanding of the battlespace across coalition partners, and 3) presentation of a unified view of the factors that influence the decisions by the coalition forces. Military planning is complex and varies in form and scope, and significant training is conducted in preparation prior to any military operation. The rationale is that after training the performance of the unit can be shared at the commander level to ensure commonality of level of preparedness across coalition partners. The same material can be used as a part of the mission planning where details of potential threats and vulnerabilities can be discussed. In [4] a detailed description of AADT is given. In this section, we will discuss how AADT DA can be used for information sharing to support training/mission planning.

The same modules that the commanders used to share information in the previous section (visual search task) can be adapted as training modules and integrated into a course to be administered to the soldiers within the coalition units. Data from individual soldiers in a unit collected from the training using the AADT can be viewed directly in the AADT DA software. This performance data collected provides the number of correct anomalies (targets and areas of interest) detected by the soldiers as progress through the module(s) is made. At its most basic, this information can be consolidated to give a score indicating correct detections, which can be reviewed and assessed for one soldier or groups of soldiers. For example, if there are soldiers within the units that have similar MOSs and carry out similar tasks, AADT DA can be used by the commanders to verify that these soldiers have a level of proficiency that will allow them to work efficiently as a team. AADT DA allows the commanders to view the information within the course(s) allowing for effortless comparisons between the performance data and ensures the integrity of data analysis. By providing this information to commanders prior to execution of a mission, gaps within preparation, whether low overall scores or low scores for a particular scenario can be addressed. Information can then be shared across the commanders safeguarding their soldiers by allowing further training or clarification of objectives. AADT DA also displays detailed statistics about keywords used in the feedback, and how these terms are used across the coalition partners and teams. In addition, the statistics show how the terms are connected to types and categories of the anomalies allowing for generation of a common lexicon and review verifying consistency in key terminology used by the coalition partners and teams. AADT DA displays detailed information about where and the sequence in which the anomalies are selected by the soldiers. This gives commanders the opportunity to confirm that the soldiers are coordinated with identifying primary threats in the operation environment. The data analyzer displays the images with the polygons and feedback created by the commanders with the corresponding highlights providing dissemination of knowledge both expert and intelligence that is useful for the task. Additionally, the data analyzer can display data from multiple soldiers at once for given image(s), which may indicate trends and patterns. This is done by using the 'Click Clustering' option shown in Figure 3. With this option, the image or images can be viewed with the marks to show where soldiers have detected potential targets or threats. Clustering of these marks can reveal common areas of vulnerabilities recognized by soldiers from the coalition forces. The clustering can also reveal vulnerabilities that soldiers from one country observes and not the other. By sharing this information a commander may be able to leverage how a partnering country's force identify possible threats and implement that strategy for the combined effort improving mission success.



Figure 3. 'Click Clustering' Option

6 – AADT and AADT DA Addressing Information Sharing Challenges

Information (knowledge) is the most important capital of current organizations and key to gaining a strategic advantage [5]. How information is obtained, stored/organized and more so how information is accessed/shared are vital considerations. In highly turbulent environments, such as military operations, information is needed even more in mission-critical situations [6]. In 2003, the US developed Blue Force Tracker to display the locations of friendly forces reducing the number of casualties. The UK enabled commanders to compile information from multiple sources allowing redirection of troops' movement to avoid evolving threats. These are examples of how individual countries have strengthened their capabilities by sharing information within their organizations [7]. Thus, sharing information "has become more significant as countries join together". There continues to be a need for automated capabilities for sharing command and control information and situational awareness information between nations [8]. The AADT and ADDT DA addresses this challenge, specifically for the identification of visual cues of potential targets and AOI for enhanced situational awareness and implications on decision making. AADT DA allows images and videos of the scenes to be viewed, reviewed with annotations, and

analyzed with expert (commanders and unit leaders) feedback from multiple sources (i.e., coalition partners). The AADT DA can be used to share information for training soldiers across coalition partners prior to missions. This gives commanders the capability of reviewing and understanding individual TTPs as well as leveraging information and perspectives from each other. It also allows the opportunity to synchronize any differences in training so that partners are on the similar pages as they equip their soldiers for engagements. The AADT DA can be used to conduct mission planning and mission reviews allowing coalition partners the ability confirm decisions for actions taken based on viewing information that is a consolidation to their individual TTPs. This also allows them to review how this information influenced the success of the mission and modify information to improve situational awareness across the coalition operation. These capabilities allow coalitions to share information supporting enhanced situational awareness that represents a more common view and assessment of the battlefield.

7 – Future work

The primary use of AADT and AADT DA has been for training but now the software is being expanded to support other areas of research. Leaders must employ both commonsense and expert knowledge when assessing threats and understanding the battlefield from visual evidence. In the future automated systems will be integrated into these coalitions. Information will then need to be shared by both leaders and automated systems. In this application, several assumptions are made including those related to security and networking. Therefore, future research that addresses the challenges such as ensuring security levels are appropriate while maximizing sharing between coalitions will provide solutions for advance interactive knowledge capture, annotating, and sharing information building on AADT and AADT DA.

8– References

- [1] “JP 5-0, Joint Operation Planning,” 2011.
- [2] J. C. Teague, “INFORMATION SHARING CHALLENGES IN A COALITION ENVIRONMENT,” 2009.
- [3] C. W. Choo, “Information culture and organizational effectiveness,” *Int. J. Inf. Manage.*, 2013.
- [4] A. J. Raglin and A. Harrison, “The Applied Anomaly Detection Tool (AADT): a platform for training Warfighters in battlefield assessment scenarios,” *2ICCRTS (21st Int. Command Control Res. Technol. Symp.*, pp. 1–10, 2016.
- [5] A. Nagendra and S. Morappakka, “Knowledge sharing barriers in global teams_ Journal of Systems and Information Technology_ Vol 15, No 3,” *Indian J. Sci. Techonology*, vol. 9, no. 45, 2016.
- [6] N. B. Jones and J. F. Mahon, “Nimble knowledge transfer in high velocity/turbulent environments,” *J. Knowl. Manag.*, 2012.
- [7] H. Trepant, M. Jansen, A. Lamaa, and A. Suddards, “Achieving information superiority Five imperatives for military transformation,” *Strategy&PWC*, pp. 1–11, 2014.
- [8] C. E. Phillips, T. C. Ting, and S. A. Demurjian, “Information sharing and security in dynamic

coalitions,” in *Proceedings of the seventh ACM symposium on Access control models and technologies - SACMAT '02*, 2002.