Understanding Distributed Communication: Visualising Action, Awareness and Intent in a Network Environment

Christopher A. Porter (Christopher.porter@dsto.defence.gov.au)
Defence Science and Technology Organisation, Edinburgh, SA 5111, Australia

Simon Dennis (Simon.dennis@gmail.com)
Department of Psychology, Ohio State University, Columbus, Oh, 43210, U.S.A.

Christopher Turley (Christopher.Turley@adelaide.edu.au)
Department of Psychology, Adelaide University, SA, 5005, Australia

Abstract

Complex communication networks have become a common facet of military, and civilian, operations in the past century. However, much of the cognitive impact of these complex networks is yet to be understood. We have endeavoured to create an effective analogue for an operational network using the Battleships™ scenario. From this experiment we have extracted several common features impacting shared situational awareness and operational intent. We have also created a prototype intent tracking software, in order to visually monitor the flow of concepts and ideas throughout the network.
Keywords: Shared situational awareness; intent; LSA; military; network enabled operations.

(R64) DANNY THINK YOU'VE GONE TOO FAR WEST
TRYING TO LOOK AT SECOND CRASH YOU SEEM TO HAVE GONE ABOUT 4 BLOCKS WEST AND 5 BLOCKS SOUTH OVER CONTINUE TAKE THE NEXT RIGHT, TURN SOUTH BOUND.
U64 THIS IS R64
(U64) ROGER NEXT RIGHT NEXT RIGHT ALLEYWAY ALLEYWAY
KING ELEMENT THEY JUST MISSED THEIR TURN TAKE NEXT AVAILABLE RIGHT UNIFORM.
U64 THIS IS R64
GOD DAMN IT STOP, GOD DAMN IT STOP.
K64 THIS IS R64

The above is an excerpt of a transcript from U.S. military communications on the 3rd of October 1993, when 18 soldiers and 100s of civilians were killed as U.S. troops attempted to capture Gen. Mohamed Farah Aidid in Mogadishu, Somalia. Two Black Hawk helicopters were destroyed in one of the worst military disasters in recent U.S. history. While many factors contributed to the outcome in Mogadishu, one issue that has consistently been highlighted is the problems with communication that occurred. Bowden (1997) describes the scenario:

"IN ORDINARY circumstances, as close to Cliff Wolcott's crash site as they were, the convoy would have just barrelled over to it, running over and shooting through anything in its path. But with the surveillance helicopters and P-3 Orion spy plane overhead, the convoy was about to illustrate how too much information can hurt soldiers on a battlefield...

The choppers and the spy plane, flown by U.S. Navy aviators, tried to steer the convoy clear of Somalian gunfire, dodging the soldiers left and right on the labyrinthine streets. It was like negotiating a maze. But the Orion pilots were handicapped. They were not allowed to communicate directly with the convoy. Their orders were to relay all communications to the Joint Operations Center (JOC) back at the beach. So when the Orion pilots said, "Turn left," that message went first to the JOC and then to the convoy. The pilots watched with frustration as the convoy drove past the place they had directed it to turn, then, getting the delayed message, turned left down the wrong street." (chapter 12, Bowden, 1997).

Lambert (1999) suggests that we can understand (military) action as "the utilization of capability to achieve intent given awareness" - the action trinity. Perhaps the most crucial capability any force possesses is the ability to communicate through language. It is by the use of language that both intent and awareness are negotiated and decisions made. Furthermore, language is the primary overt behaviour which we can analyze in our attempt to understand the processes that govern both the decision making of the individual and the emergent decision making of groups embedded in communication networks such as that which operated during the Black Hawk Down episode. However, our ability to understand language is currently limited by our inability to formally model natural language. A formal model of the content of natural language would allow us to understand questions such as:

1. How is intent distributed within the network at any given moment?
2. How is awareness distributed within the network?
3. What factors (tactical, moral etc) led to a particular action? How is decision making distributed across the network? How is action coordinated?
4. When errors arise what was their source and how are they propagated through the network?
6. How do factors such as communication channel impact communication and overall performance?

The Battleships Experiment

Previous work on analysing distributed decision making has tended to focus on the action sequences that groups produce. Intent, awareness and decision processes are then inferred from these sequences. This approach tends to lead to over simplistic task environments that permit direct inference of psychological processes, but that lack many of the dimensions that govern performance in real settings. Alternatively, reasonably faithful simulations are employed at the expense of being able to infer in any detail the relevant thought processes (Dekker, 2006; Knott et.al. 2006). In this line of work, we have focused on formal modelling of language in reasonably simple situations, in preparation for developing a more comprehensive understanding of real world network decision making.

Design and Procedure

The current study used a Battleship ™ like scenario, adapted from the Milton Bradley board game (1931) to investigate the differences between three communication conditions (1) Voice (2) Text and (3) No communication. Distributed teams of three participants for the voice and text conditions were positioned in separate rooms and linked via a network to a server. In the no communication condition, a single player was used to simulate complete intent and situational awareness sharing. Each participant was required to complete three rounds. A round ended once all ships on the game board were hit and destroyed. The multiplayer conditions required participants to communicate their shot intention and results to one another in order to develop a shared situational awareness during the game. The Battleships game was used because it is widely known and easily learnt. Therefore, much of the difficulty in training participants to use the game effectively was circumvented, reducing learning effects across rounds.

Each team, or individual in the case of the single player condition, began each round with 1000 points, and the game was scored decrementally at the end of each turn. In order to promote speedy ship destruction the points lost were a function of number of ships remaining on the board and the number of active-squares on the game board, where an active-square is any square occupied by a ship. There are six randomly placed ships in total, with the following square distribution: 1 x Carrier (5 squares long), 1 x Battleship (4 squares long), 1 x Submarine (3 squares long), 1 x Destroyer (3 squares long) and 2 x Cruiser (2 squares each), amounting to 19 active squares or 19% of the board. Ships can only be placed vertically and horizontally, therefore there were no diagonally orientated ships.

Figure 1 shows the game board as it appeared to participants. The left mouse button was used to register a shot intention, i.e. where the player wished to fire (denoted by a hollow green circle). Shot intentions could be changed to another location by left clicking on the new target square within the turn time limit. The right mouse button was used to mark the game board. There were three different mark types: right clicking on a selected square

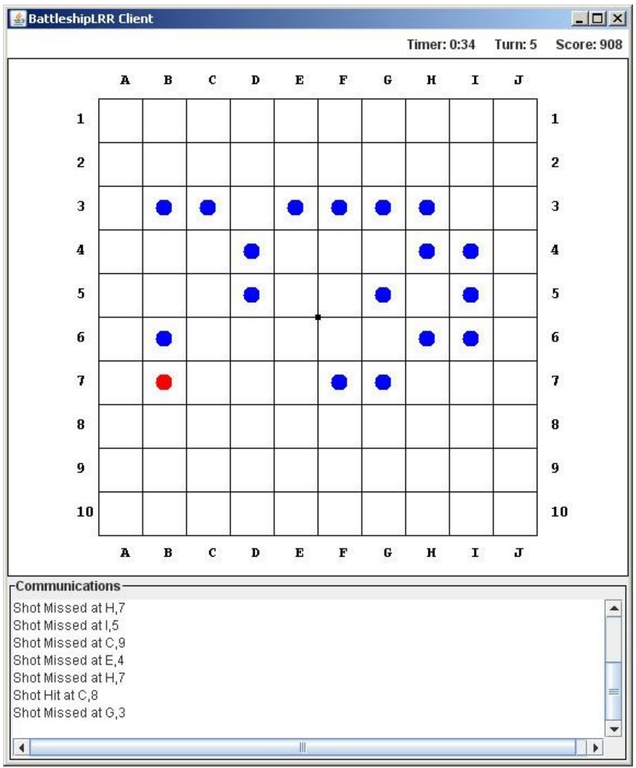

Figure 1: Battleships Game Board

once would place a yellow circle in the square; denoting an area of interest. For example, in the multiplayer case the shot intentions of other team members. A second right click on the same square changed the yellow circle into a blue circle; denoting a miss. A third right-mouse click on the same square changed the blue circle into a red circle; denoting a hit. After each turn elapsed players received textual feedback on their shots and were required to visually update their own boards. In this way, we were able to gather a detailed understanding of their situation awareness. Upon destruction of a ship players were informed of the ship type and size, as per the original game.

Results and Discussion

In this section, we draw attention to three main variables. These include the analysis of score, which might be thought of as an analogue of overall mission effectiveness and the communicated intent changes and communicated confirmations, which capture the notions of intent and situation awareness components of the action trinity proposed by Lamberts (1999).

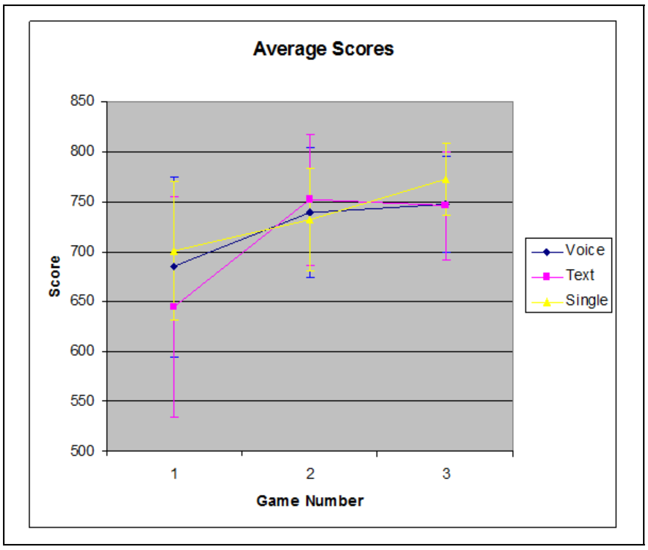

Figure 2: Average Scores

Score. Figure 2 shows the graph of the means of the scores across the three games for the communication conditions. There was a significant increase in scores as a function of game demonstrating a learning effect. There are two main aspects that participants could be acquiring as they become more practiced, the rules of the game and the methods for communicating. The fact that the single players show only modest increases across games suggests that scores are improving primarily because participants are developing more efficient methods of communication.

Surprisingly, the between communication conditions effect was not statistically significant indicating that communication channel did not impact overall performance. It may have been that the turn duration of one minute was too long to show a substantive advantage for the high bandwidth voice channel. As we will see in subsequent sections, however, the style of communication was affected by channel type.

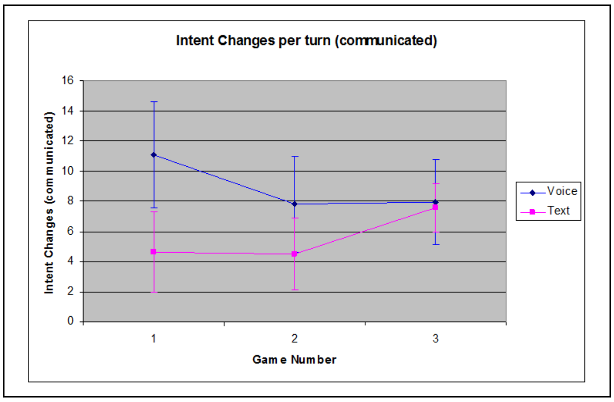

Figure 3: Communicated Intent Changes

Intent Changes. A second variable that did demonstrate a substantive difference between communication channels was the number of communicated intent changes (see Figure 3). The primary cause of a change of intent is the introduction of novel information to the player from other members of the team, such that they alter their target. A common example of this is the change of targeting to sink a newly located ship.

At first participants in the voice condition made more communicated intent changes than participants in the text condition. This may be due to a higher cognitive load for participants in the text condition learning to communicate, via reading messages and typing to respond, than their counterparts in the voice condition. In turn this may lead to fewer changes to begin with. However, as participants became more proficient the number of intent changes seemed to stabilize on approximately eight irrespective of the channel used. This pattern of results suggests that there may be an optimal number of intent changes and that ultimately either channel arrives at this number.

Confirmations. The final issue that we will highlight focuses on participant's efforts to maintain situation awareness. In particular, we look at their communications that involve confirmations of the status of squares on the game board. We divide these up into proximal and distal-confirmations in order to differentiate those that involve a confirmation of a communication from the current turn from those that involve communication from previous turns (see Figures 4 and 5).

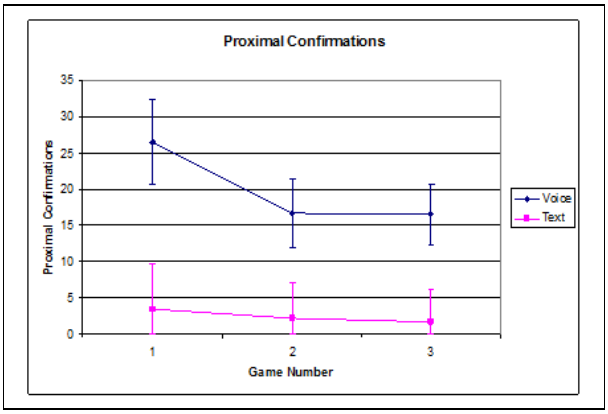

Figure 4: Proximal Confirmations

The pattern of results is striking. While distal confirmations show little difference between communication channels there is a distinct advantage in the text channel when proximal confirmations are considered. Note the need to issue a distal confirmation in either channel relies on the participant's memory and so the requirements are similar. In the proximal channel, however, it may be that the participants have some difficulty identifying what has been said thus generating a

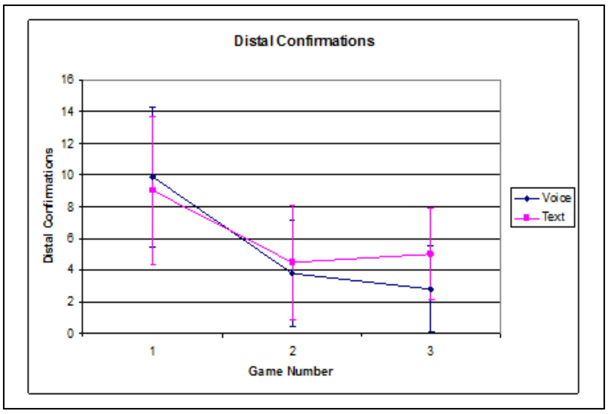

Figure 5: Distal Confirmations

request for confirmation that would not be necessary in text where there is little distortion. In addition, if a participant has forgotten a recent communication in the text mode it is likely that it will be still visible and so no confirmation will be required. However, over a longer period of time the communication will have scrolled out of the window, and thus requesting confirmation requires less cognitive load than finding the previous communication.

Visualizing Network Communications

The second objective of the project was to develop software that is capable of visualising how distributed communications are being utilized to maintain situation awareness and realize intent. Most network visualisation software focuses on identifying the amount of traffic occurring at a given location at a given time. In order to provide insight into the utility of communication in an operational setting, however, we need to understand not just how much traffic is occurring but what the content of that traffic is. Our objectives are to be able to provide a tool that can identify when and where communication difficulties that have lead to an adverse result have occurred and ultimately to anticipate and automatically circumvent any bottlenecks that may lead to problems. For these purposes, we require formal models of message content.

To achieve this we have used Latent Semantic Analysis (LSA, Landauer & Dumais, 1997; Landauer, McNamara, Dennis, & Kintsch, 2007). LSA is a statistical natural language processing technique that takes a large corpus of text and produces vectors representing words. A large occurrence matrix indicating which words appear in which documents of the corpus is analysed using singular value decomposition. By retaining a relatively small number of factors (typically about 300) noise is filtered from the large raw matrix and vectors of words that have similar semantic content are located in similar areas within a Euclidean space. These word vectors can then be added to construct vectors representing individual messages.

Message similarity is determined by taking the cosine between these vectors. The method has been used in a number of applications including essay grading, information retrieval and automatic discussion group monitoring (Landauer et al., 2007).

The key reason for employing a method like LSA is to be able to capture semantic content regardless of the words that are used to express that content. Messages can typically be expressed in many different ways and any method that relies on exactly matching words or on natural language engineering approaches is likely to be brittle and unreliable. LSA circumvents these problems.

Figure 6 shows a prototype interface that we have created. Each of the boxes represents a node in the communication network and the arcs indicate the communication channels between them. In the Battleships experiments this graph is simple consisting of three nodes and broadcast communications. The textbox below the graph allows the input of a query. This query is simply a piece of text that expresses the content that we are interested in at this particular time. This might be an intent statement such as: “let’s take out the cruiser in the bottom right hand corner” or it may be an element of situation awareness such as “hit at D4”.

The colouring of the nodes indicates the extent to which the communication of the node in the current turn matches the query content. In this case, player 1 (top left hand node) has just said ”d4 = Hit” and player 2 (bottom node) marked D4 as a hit. The LSA cosine to the query is high in both these cases indicating a match and so these nodes have a redder hue. To observe how content distributes across the network the transcript can be played turn by turn, and the communications visually seen.

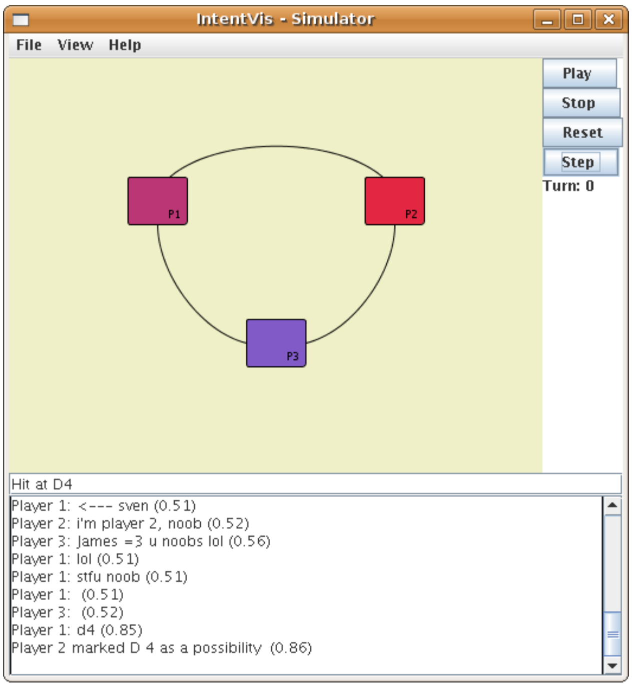

Figure 6: Intent Visualisation Software

While this is a very simple example, it illustrates two important aspects of the interface. Firstly, it is not necessary for the communication to be expressed in exactly the same terms for the method to identify that the message is relevant. Secondly, we are able to incorporate actions and verbal communications into the same interface thus allowing us to utilize all available evidence to infer the situation awareness of the participant.

Original trials of this method were done with a complete natural language corpus, however later trials have used a customised corpus in order to increase the fidelity of the extraction. Specifically square locations can now be regarded, by LSA, as being closer, more related, or further away, less related. This location based information allows us to not only visualise the interactions for a specific query, but also to see some of the inferences which the players may have made as well.

Of course, the broadcast nature of the current experiment is not ideal for demonstrating the power of the method. We have started collection of new data in which participants are arranged in a chain and are only able to communicate with the adjacent participants. Such data should provide a more meaningful test of the interface as it should then be possible to track communications as they move from one end of the chain to the other. A key test is how well we are able to maintain a lock on those communications as each member of the chain chooses their own way of expressing the content.

Conclusions

The Battleships scenario has proven a useful framework to begin understanding a number of issues in distributed communication its impact on of mission effectiveness, intent and situation awareness. We have seen that the Battleships scenario can indeed extract useful information, especially with regard to the participant's situational awareness, and their stated intents.

We have found that participants show different styles of communication and different patterns within each communication condition, but these conditions prove statistically similar in overall effectiveness. The study has also yielded significant information for the analysis of network communication in the form of a large corpus with which to analyse the textual data, and also investigate the efficacy of tracking intent and situational awareness across these networks during the trial.

We have found that it is possible to track situational awareness and intent across the network, and that the methodology to execute this is not overly intensive, which may allow for future real-time concept tracking. A further extension on this work is currently being conducted in order to investigate the possibility of tracking intent and situational awareness across a chain type communications structure. In addition this extension seeks to rectify some of the possible learning influences present in the game, so as to provide further fidelity.

More development is continuing on the prototype network visualisation software to expand this into a suite capable of analysing arbitrary communications. Future work may include the analysis of real-world communications transcripts, such as that available from the Mogadishu episode discussed in the introduction. More data collection may be undertaken in order to use a more complex, and faithful, analogue for further comparison and expansion.

There is little doubt that the effectiveness of complex communications networks will continue to be an important factor in the pursuit of military, and civilian, objectives in the future. Our empirical studies and the novel visualisation software that we have developed offer a unique perspective on the analysis and effectiveness of network enabled operations in the future.

Acknowledgements

We would like to thank the Defence Science and Technology Organisation (Australia) for their continued support of the project, and the Psychology Department of Adelaide University.

References

Bowden, M. (1997). *Black Hawk Down: An american war story*. (http://inquirer.philly.com/packages/somalia/)

Dekker, A.H.(2006) *Revisiting "SCUDHunt" and the Human Dimension of NCW.* Transcripts of the 11TH ICCRTS.

Knott, B.A., Bolia, R.S., Nelson, W.T., Galster, S.M. (2006). *Effects of Collaboration Technology on the Performance of Tactical Air Battle Management Teams.* Transcripts of the 11TH ICCRTS.

Landauer, T. K., & Dumais, S. T. (1997). *A solution to plato's problem: The latent semantic analysis theory of the acquisition, induction, and representation of knowledge*. Psychological Review, 104, 211-240.

Landauer, T. K., McNamara, D. S., Dennis, S., & Kintsch, W. (Eds.). (2007). *Handbook of Latent Semantic Analysis.* Mahwah, New Jersey: Lawrence Erlbaum Associates.