13th ICCRTS: C2 for Complex Endeavors

Networked Collaborative Intelligence Assessment

Topic
Collaborative Technologies for Network-Centric Operations

Authors
Matthew R. Risser, Ph.D.
Harvey S. Smallman, Ph.D.

Point of Contact
Matthew R. Risser, Ph.D.

Name of Organization
Pacific Science & Engineering Group

Complete Address
Pacific Science & Engineering Group
9180 Brown Deer Rd
San Diego, CA 92121

Telephone
858-535-1661

Email
risser@pacific-science.com

Networked Collaborative Intelligence Assessment

Abstract

There is a recognized need to improve the exchange of information among the intelligence community; however, the barriers to effective sharing are high. One barrier is the lack of dedicated collaborative intelligence assessment tools and concepts of operation for their usage. As part of the Office of Naval Research (ONR) Collaborative Knowledge Interoperability (CKI) program, we are developing a networked collaborative intelligence analysis tool called JIGSAW (Joint Intelligence Graphical Situation Awareness Web). JIGSAW contains a shared graphical workspace for analysts to post and share their assessments to structure and facilitate their collaboration. This paper examined alternate JIGSAW collaborative architectures, who can see what and who can talk to whom, to determine their effect on collaborative assessment. In a four week *in vivo* evaluation, a web-based JIGSAW prototype was used by a Naval Postgraduate School (NPS) decision support class to work ten quasi-topical strategic analysis problems. Two design factors were manipulated: (1) the presence or absence of others' assessments in the visual workspace and (2) a centralized or decentralized collaborative discussion architecture mediating analyst discussions. Collaborative assessments were tracked over time and measured by a novel dispersion metric. The results revealed that the nature of the discussion architecture affects how quickly the assessment of analysts converges over time. Further, there was a noticeable trend to shift assessments towards confirming the hypotheses in the problems, suggestive of group confirmation bias. The results of the study can be used to inform the design of networked collaborative analysis systems to address decision biases and facilitate effective information exchange.

Keywords: intelligence analysis, collaboration, assessment, decision bias, collaboration metrics

Introduction

As the 21st Century begins to unfold, the US military continues to transform itself in the face of powerful global forces that impact its missions and the intelligence required to successfully conduct them. First, missions often have been prosecuted against dispersed and rapidly mutating foes - different enemies, indeed, from the monolithic, slower moving and more predictable Cold War opponents that much of the intelligence infra-structure was set up to analyze. Second, the computing and internet revolution that has changed the information landscape continues to quicken and intensify. Networked technologies offer opportunities to distribute complex cognitive work activities to dispersed teams of information warriors and analysts to combat these new challenges of asymmetric warfare. The US military has recognized these possibilities and embraced these networked technologies, using a net-centric model. In this model, C2 and supporting intelligence will be conducted by diverse and distributed collaborating teams of operators and analysts. Vast new networked infrastructures are being developed by the

services along with the tools and technologies to ride on these networks to support this net-centric warfare model. However, there are many unanswered questions about how to design these tools and develop associated concepts of operation for their usage to ensure that their human users can effectively collaborate and perform the complex cognitive analysis tasks required of them in this future model. This paper reports an evaluation of novel cognitive designs to support the specific problems of networked collaborative intelligence assessment in that net-centric model.

There are many reasons to support improved sharing in intelligence analysis. Some reasons stem from external pressures. For example, the 9/11 commission report highlighted the need for intelligence sharing – specifically inter-agency sharing – to "connect the dots" to prevent future intelligence failures (The 9/11 Commission, 2004). At the individual analyst level, the few systematic cognitive task analyses that have been conducted with analysts have reported a fairly consistent set of issues related to problems in the areas of business process, technology, and the cognitive challenges of integrating information and making unbiased assessments (Patterson, et al., 2001; Hutchins, et al., 2004; Smallman, 2008). Other reasons stem from the need for shared understanding among distributed analyst teams (e.g., Warner, Letsky & Cowen, 2005; St. John, et al., 2005). For an excellent older, but still very relevant discussion of some of the contradictions inherent in trying to make sharing and more effective collaboration work with classified material work in a strict hierarchical organization, see Hall (1998).

How should sharing be supported in intelligence analysis? What should the business processes and concepts of operation be? How should future networked collaborative tools be designed to support users in getting the information they require, integrating it, and then drawing unbiased assessments about it? For answers to these questions, designers and developers can draw on work conducted in a variety of different disciplines, including computer supported collaborative work (CSCW), the psychology of decision making (e.g., Hogarth, 1987), and the cognitive science of team collaboration (see Letsky, Warner, et al., 2007). As part of the ONR-sponsored Collaborative Knowledge Interoperability (CKI) program, the JIGSAW project (Joint Intelligence Graphical Situation Awareness Web) is developing a networked collaborative intelligence assessment system that integrates relevant concepts from these three areas with several new innovations to address these decision making challenges. JIGSAW is web-based system intended to facilitate the exchange of information and reduce biases among disparate and unfamiliar intelligence analysts (for review, see Smallman, 2008). The collaborative nature of JIGSAW is facilitated primarily through two design features. First, there is a shared visual workspace or intelligence landscape (I-Scape) that is currently instantiated. Second, there is a putative discussion mechanism under development whereby analysts can share and evaluate information from unknown but knowledgeable sources. Both features are intended to encourage collaboration and reduce the "usual suspects" phenomenon where analysts tend to exchange information with only those analysts who are familiar.

The I-Scape feature is specifically intended to address the interpretation of information in the collaboration process. It was designed to reduce biases by visually representing information on two dimensions: credibility and support. For an analyst, this makes

conflicting evidence salient through a graphical layout because it stores the interpretation with the information, minimizes reliance on memory, and increases the availability of information. In a set of controlled studies, Cook and Smallman (2007, 2008) examined the effect of visual debiasing as compared to the textual representation of the same information with trainee naval analyst participants. Further, they investigated the debiasing effect of having intelligence assessed by analysts other than the primary analyst tasked with working a problem. They found a significant bias to focus on and share confirming evidence – the classic confirmation bias (for review, see Nickerson, 1998). Further, the highest ranked evidence was also the most supportive evidence. However, Cook and Smallman did not find a significant debiasing effect of others' assessment. One limitation of their study was that the experiment was not truly collaborative in nature; thus, requires additional study to determine if the effects persist when used in an asynchronous collaborative environment. Therefore, one motivation of the present study was to see if the confirmation bias was evident when JIGSAW was used in truly collaborative, realistic evaluations, by manipulating the visual presence or absence of others' assessments on the I-Scape display (e.g. visually displaying other's assessments of information or only your own). The study sought to determine whether or not the ability to see graphical representations of others' evidence had any effect on the interpretation of ones' own evidence when working intelligence problems.

As part of the collaborative process, the technology needs to facilitate the shared understanding of information which requires some method to support discourse. For example, chat rooms are a tool in ubiquitous use by the military for information sharing (e.g., Heacox et al., 2005). However, chat rooms have numerous disadvantages for users - they do not store or represent discourse well and were mostly intended for synchronous communication only. For the purposes of intelligence discourse, an asynchronous means for communication is required because intelligence problems tend to be longer-term in nature and the future net-centric model will see their use in asynchronous collaboration. Thus, in this study, a discussion forum was used to mediate discourse because of its permanence. Similar to the decision bias problems noted above, open discussion forums can have implicit biases due to social influences of contributors. Therefore, the present study manipulated discussion architectures to better understand how consensus of assessment might change over time as JIGSAW is used to solve real problems. Discussions were routed through the primary analyst responsible for working the problem. In the "centralized" architecture (one-to-many), the primary analyst could see all discussion threads with each contributor whereas contributors could only see their own discussion with primary analyst. On the other hand, the "decentralized" architecture was an open forum moderated by the primary analyst. It differed from the centralized architecture in that the contributors could see all of the discussion threads from other contributing analysts. Therefore, it was open, but not a true peer-to-peer topology because of the role of the moderator. The centralized discussion/I-Post absent condition combination is a rough proxy for today's analysis and the decentralized discussion/I-Post present condition combination is a proxy for a future JIGSAW analysis model. To obtain as much information as possible about JIGSAW's realistic usage but still manipulate these two variables parametrically, the experiment was conducted as an *in vivo* four-week study of JIGSAW used collaboratively to work real problems by a class of Naval Postgraduate School students, several of whom possessed analysis experience.

Method

In general, the 4-week experiment required the collaborative analysis of ten problems. Participants role-played analysts proceeding through the following sequence of tasks: evaluate 10 hypotheses, find intelligence related to the hypotheses, participate in discussion, and then re-evaluate the intelligence. The Concept of Operations (CONOPs) for JIGSAW envisions two roles for collaboration, the primary analyst (hypothesis owner) and other contributing analysts. In JIGSAW, the pieces of intelligence are posted as visual objects or Intelligence Posts (I-Posts) within a shared visual assessment space called an Intelligence Landscape (I-Scape). The discussion forums progressed within two collaborative architectures: centralized (owner-to-contributors, or one-to-many) and decentralized (moderated open forum). The I-Scape presented I-Posts in two arrangements: the presence (view all I-Posts) or absence (only view own I-Posts) of others' I-Posts. Therefore, the collaborative evaluation of intelligence was investigated on the basis of both the arrangement of I-Posts on an I-Scape and the scope of discourse as a function of discussion architecture.

Participants

The study included twenty participants from a Naval Postgraduate School (NPS) Decision Support Systems (DSS) class (19 students/military officers and one instructor). There were no restrictions on gender, age, or ethnicity for inclusion into the study. Collectively, the participants represented the Navy, USMC, and USCG and 25% had intelligence experience. The average age was 37, and participants had an average of 13 years of military experience.

Design

The study used a 2 x 2 x 2 mixed design that included I-Scape arrangement (presence vs. absence of others' I-Posts on the display) and time (initial vs. final assessment of intelligence) as within-subject variables, and discussion architecture (centralized vs. decentralized) as the between-subject variable.

The I-Scape is a graphic analysis space populated with collaborators' intelligence assessments. The 2D Cartesian nature of the I-Scape afforded development of a novel metric of the spatial dispersion of I-Posts on the I-Scape. It was calculated by taking the square root of the sum of the variance for x and y I-Post coordinates on the I-Scape. A small dispersion might reflect spatial "clumping" of assessments, and hence an indicator of group consensus on the hypothesis. A large dispersion, on the other hand, would be indicative of more heterogeneous assessments and hence less group consensus. This dispersion metric became the primary dependent variable. The meaning of the I-Scape's x and y axes are described in more detail below and shown in Figure 2.

Additional dependent variables included I-Post position on both support and credibility axes, frequency of I-Post assessments, net distance of I-Post movement, total distance I-Post movement, and number of discussion entries.

Stimuli

Hypotheses. The primary stimuli for the experiment were ten fictitious analysis problems which were in the format of ten separate binary hypotheses and created based on the following criteria:

- different topics that were current and of general interest
- referenced by multiple open sources
- binary (true/false)
- balanced (likely equal evidence pro and con)
- implications for government response (military or non-combatant)

The 10 Hypotheses listed below were generated from open source news websites such as CNN, circa 2006.

1. Turkey will be admitted into the European Union starting with the next budget cycle in 2014.
2. Iran is pursuing the development of nuclear weapons
3. By the end of 2008, the U.S. will be prepared to effectively respond to any Avian flu outbreak
4. In order for the new Iraqi government troops to effectively maintain control by themselves, at least 75,000 US troops must be stationed in Iraq until 2011.
5. There will be another significant terrorist attack on U.S. soil within the next 2 to 3 years
6. Osama bin Laden will be captured by the end of 2006
7. There will be another significant natural disaster in the U.S., on the same scale of devastation as Hurricane Katrina, within the next 5 years
8. In response to the Kremlin's crack-down on Muslims, Al-Qaeda will begin terrorist activity in Russia within the next 5 years
9. Korea's rogue activities will result in a major armed conflict on the Korean Peninsula within a decade
10. Thailand will be the next country targeted for attack by the Asian Abu-Sayef terrorist group

Tools. A stripped-down and simplified web-based experimental version of JIGSAW was created for this collaboration study to focus on the design features of interest which consisted of three components: a hypothesis manager, a JIGSAW interface (I-Scape with I-Posts), and a discussion forum. Each hypothesis had its own page and each had a JIGSAW interface and discussion forum. The components were integrated and hosted on a secure web server accessed remotely by the participants.

Hypothesis Manager. The hypothesis manager was an experimental, stripped-down version of JIGSAW's RFI manager and maintained records of all activity for each hypothesis in a tabular display. This table allowed participants to see progress on each of the individual hypotheses and provided links to each of the 10 hypothesis pages. See Figure 1 for a screenshot of the Hypothesis Manager.

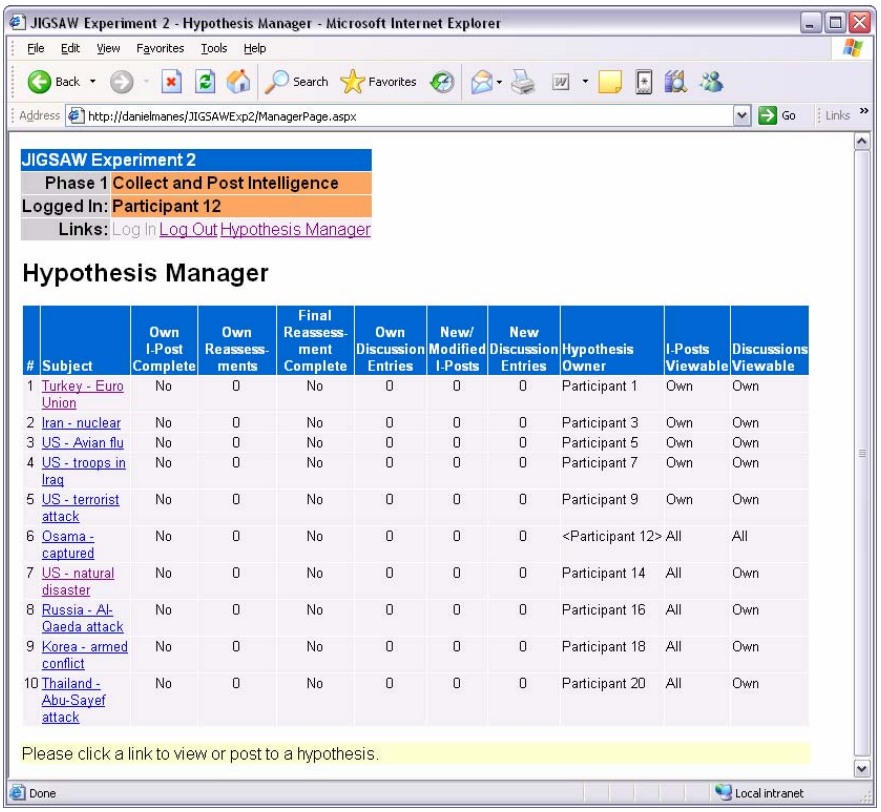

#	Subject	Own I-Post Complete	Own Reassess-ments	Final Reassess-ment Complete	Own Discussion Entries	New/ Modified I-Posts	New Discussion Entries	Hypothesis Owner	I-Posts Viewable	Discussions Viewable
1	Turkey - Euro Union	No	0	No	0	0	0	Participant 1	Own	Own
2	Iran - nuclear	No	0	No	0	0	0	Participant 3	Own	Own
3	US - Avian flu	No	0	No	0	0	0	Participant 5	Own	Own
4	US - troops in Iraq	No	0	No	0	0	0	Participant 7	Own	Own
5	US - terrorist attack	No	0	No	0	0	0	Participant 9	Own	Own
6	Osama - captured	No	0	No	0	0	0	<Participant 12>	All	All
7	US - natural disaster	No	0	No	0	0	0	Participant 14	All	Own
8	Russia - Al-Qaeda attack	No	0	No	0	0	0	Participant 16	All	Own
9	Korea - armed conflict	No	0	No	0	0	0	Participant 18	All	Own
10	Thailand - Abu-Sayef attack	No	0	No	0	0	0	Participant 20	All	Own

Figure 1: Hypothesis Manager interface that provided experiment progress and status.

Intelligence Posts (I-Posts). Each I-Post used a web-based form to capture the author, time, date, source, message, and ratings. I-Posts were rated on two factors: *credibility* and *degree of support* on two 7-point scales. Ratings determined the placement of the I-Post on the I-Scape in an x-y coordinate space. Participants also collected and provided their own intelligence related to each hypothesis. Intelligence could be in the form of any open source information (e.g., referenced text, images, documents, tables, web pages).

Intelligence Landscape (I-Scape). The I-Scape interface supported the placement of an I-Post and editing its position on both 7-point credibility and support axes. Each I-Post was displayed in the I-Scape and the information was automatically captured in the discussion forum. Thus, each I-Post within an I-Scape was hyperlinked to its respective hypothesis-based forum. See Figure 2 for a screenshot of the I-Scape and I-Posts.

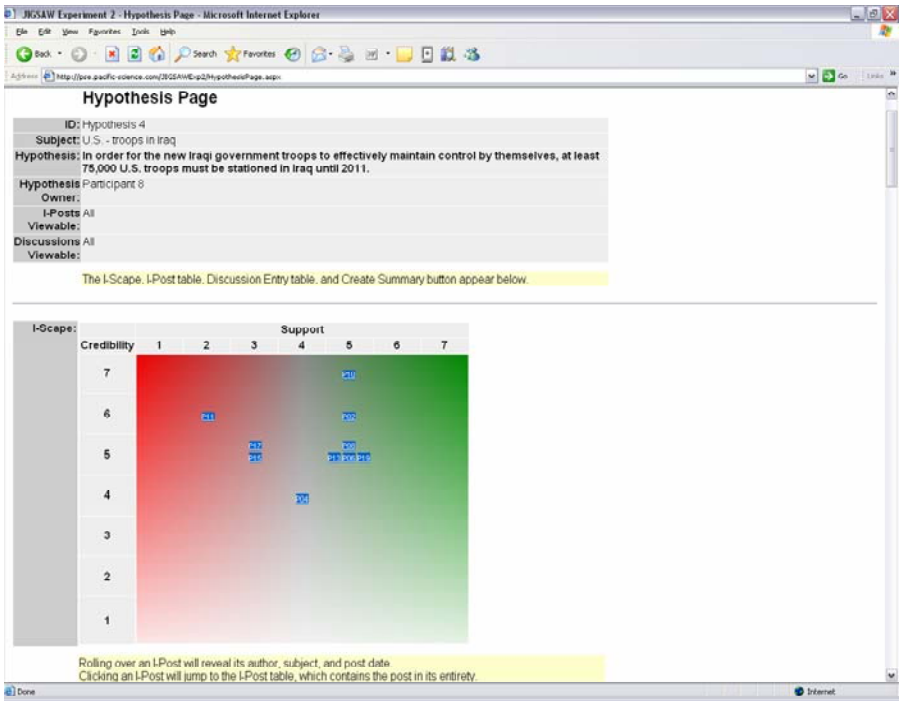

Figure 2: Sample I-Scape and I-Posts interface for collaborating on each hypothesis.

Discussion forum. The discussion forum was used to support hypothesis collaboration via two discussion architectures: centralized (one-to-one discussions between the hypothesis owner and contributing analysts) and decentralized (moderated open forum discussions between all members). The centralized discussion architecture supports the traditional flow of communication where the hypothesis owner communicates one-to-one with contributing analysts. On the other hand, the decentralized architecture allowed contributing analysts to communicate with each other with the hypothesis owner acting as the moderator. Figure 3 represents the two discussion architectures.

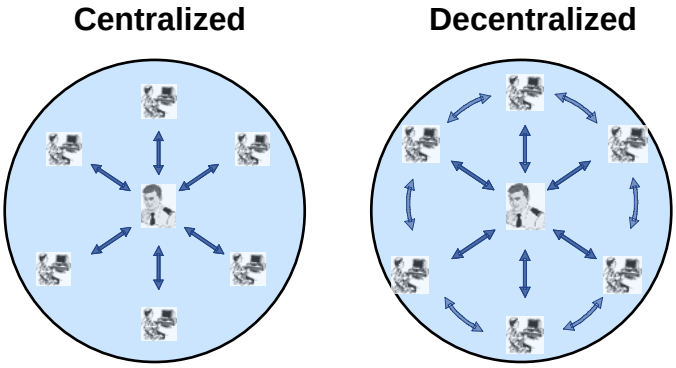

Figure 3: Schematic of the Centralized (left) and Decentralized (right) discussion architectures manipulated in the experiment.

The discussion forum was organized by hypothesis where each piece of intelligence was posted as a new message within its respective hypothesis-forum. See Figure 4 for a screenshot of the discussion forum.

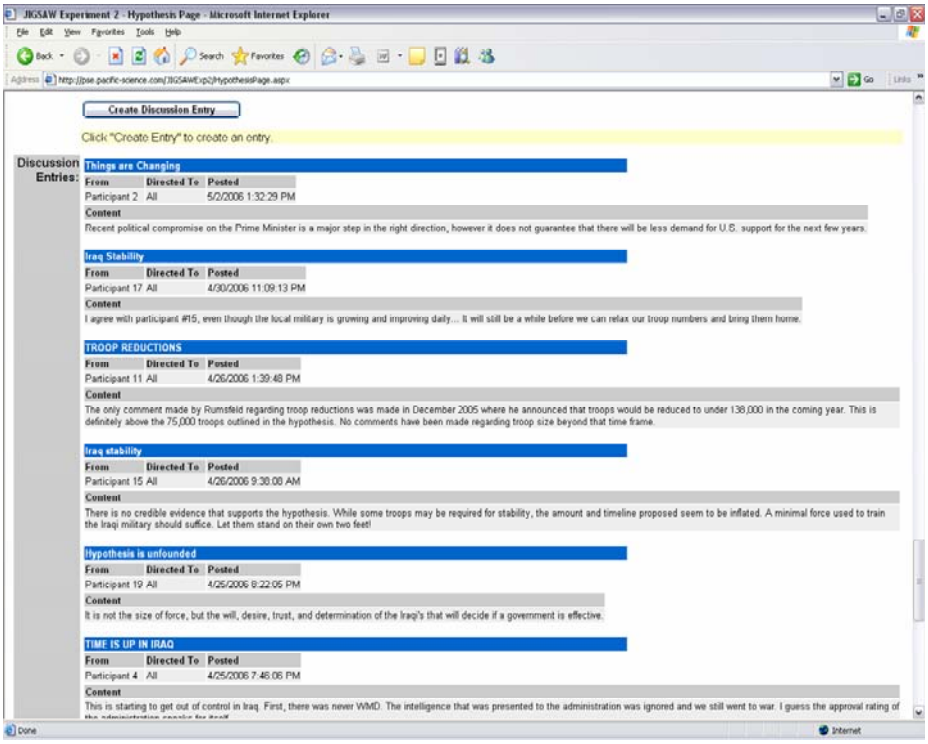

Figure 4: Hypothesis page with discussion forum entries submitted by participants role-playing analysts collaborating and working problems.

Procedure

The 20 participants were given a goal to work as a team of analysts using the new JIGSAW tool and were assigned to one of two groups - 10 in the centralized and 10 in the decentralized discussion group. Participants worked 10 fictitious open source and unclassified hypotheses over four weeks for approximately 15 min/day or five hours total for the experiment as a class assignment. Each participant was made owner of one hypothesis and a contributor to all 10 hypotheses. The same 10 hypotheses were assigned to both centralized and decentralized discussion groups. Both I-Post and discussion activities could be performed asynchronously among team members. All forums were periodically monitored by the experimenter to ensure the experiment was progressing as it should and that there was no unwarranted conversation. The details of the procedure can be summarized into four phases: (1) intelligence collection and posting, (2) discussion, (3) re-evaluation, and (4) final reporting. See Figure 5 for a schematic diagram of the experimental procedure with I-Scape arrangement and discussion architecture.

Phase I: Intelligence Collection and Posting (12 days). Participants were given 12 days to collect and post intelligence as an I-Post in the appropriate I-Scape. They logged into JIGSAW and were presented with the list of 10 hypotheses for their review. Participants collected one piece of intelligence for each hypothesis, including their own. Thus, each was responsible for collecting 10 pieces of intelligence and completing the I-Post form. This activity created an I-Post and a new discussion forum message for the respective hypothesis.

Phase II: Discuss Hypotheses and Intelligence (12 days). After all intelligence was posted for each hypothesis, the hypothesis owner and contributors were provided an opportunity to collaborate with one another according to the assigned discussion architectures. Participants were provided 12 days to discuss the intelligence and hypotheses. Recall that in the centralized architecture, each contributor was only able to discuss one-to-one with the hypothesis owner. Alternatively, the decentralized architecture allowed participants to communicate with other contributors as well as the owner. Server-assigned permissions controlled what participants were allowed to see (i.e., both I-Posts on the I-Scape and discussions in the forum). Participants were able to log into JIGSAW, review the I-Scapes, and change the position of their own I-Post at any time; however, they were not able to edit the content of their I-Posts or discussion messages once posted.

For each participant, five of the I-Scapes were presented with others' I-Posts (presence) while the other five were presented without others' I-Posts (absence). In other words, for half of the hypotheses, all participants' I-Posts were displayed and the other half of hypotheses only displayed one I-Post, their own. The hypotheses were randomly assigned to the presence and absence conditions for each participant.

Phase III: Review and Reassess Own I-Posts (3 Days). After discussions were completed, three additional days were provided for the hypothesis owner to review all intelligence posts and conclude any communications with contributors. After communications were completed, contributors were provided with an opportunity to review and reassess their final I-Post positions on the I-Scapes (using the method above).

Phase IV: Provide Final Response (3 Days). At the conclusion of the study, the hypothesis owners logged into JIGSAW and provided a response to explain only their own hypothesis (to support or refute).

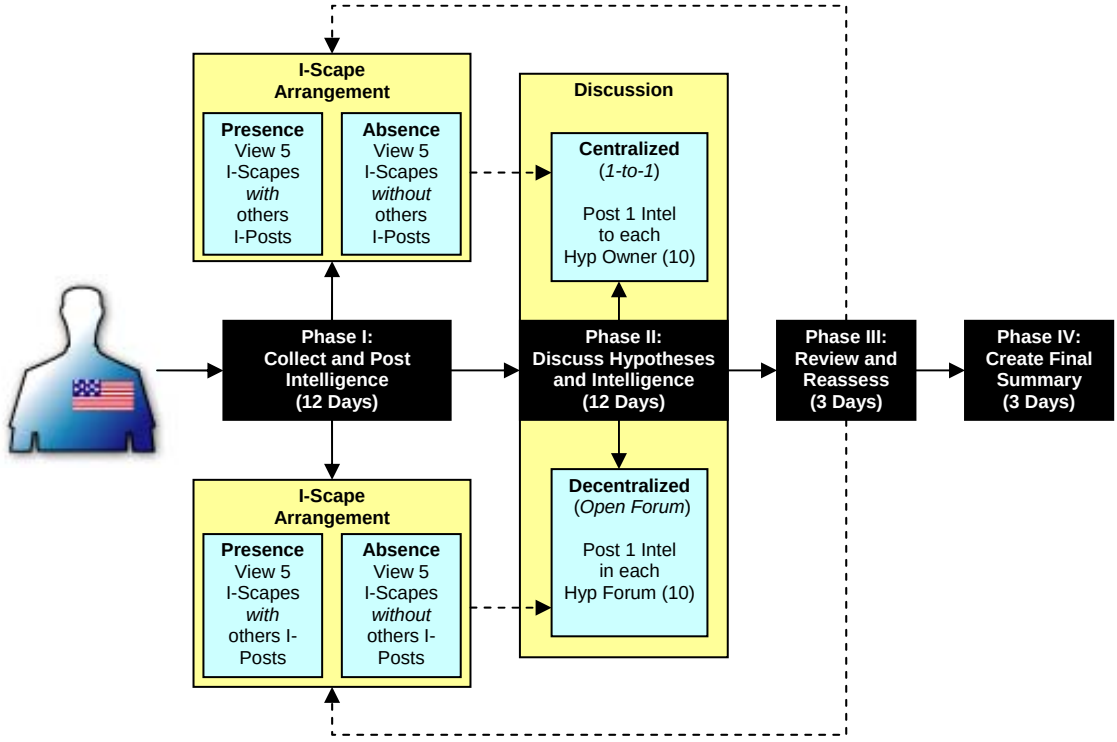

Figure 5: Illustration of the timing of the different phases of the experimental procedure with the I-Scape arrangement and discussion architecture manipulations.

Analysis and Metrics

A repeated-measures ANOVA was used to analyze discussion architecture, I-Scape arrangement, and time as measured by I-Scape dispersion, I-Post position on both support and credibility axes, frequency of I-Post assessments, net I-Post distance, total I-Post distance, and number of discussion entries. A correlational analysis was performed to determine if there was a relationship between amount of discussion and I-Post movement as measured by dispersion or distance.

The I-Scape axes of Degree of Support and Source Credibility were rated on a 7-point scale and provided the Cartesian coordinates to calculate the I-Post means, movement distance, and Cartesian measurement of dispersion, expressed as:

$$\sqrt{x_\sigma^2 + y_\sigma^2}$$

This measure of dispersion captures both axes in a single metric. However, as a scalar quantity, it does not account for the direction of the relationship. Figures 6 and 7 graphically depict the difference in the absolute measure of dispersion between a decentralized and centralized discussion architecture given the same means for each axis.

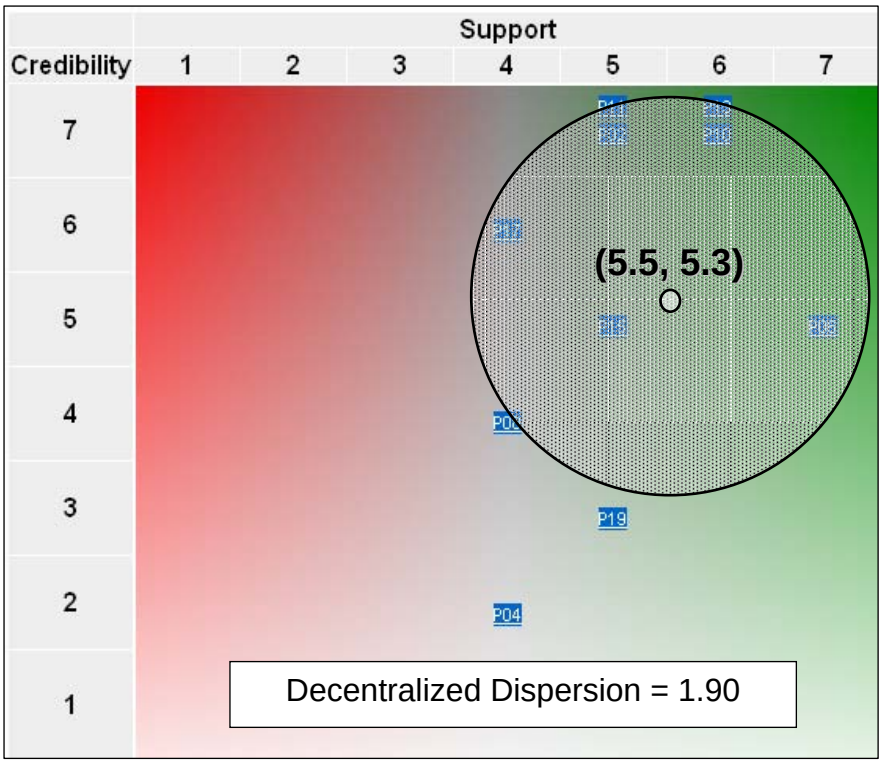

Figure 6: Decentralized discussion with a dispersion of 1.90

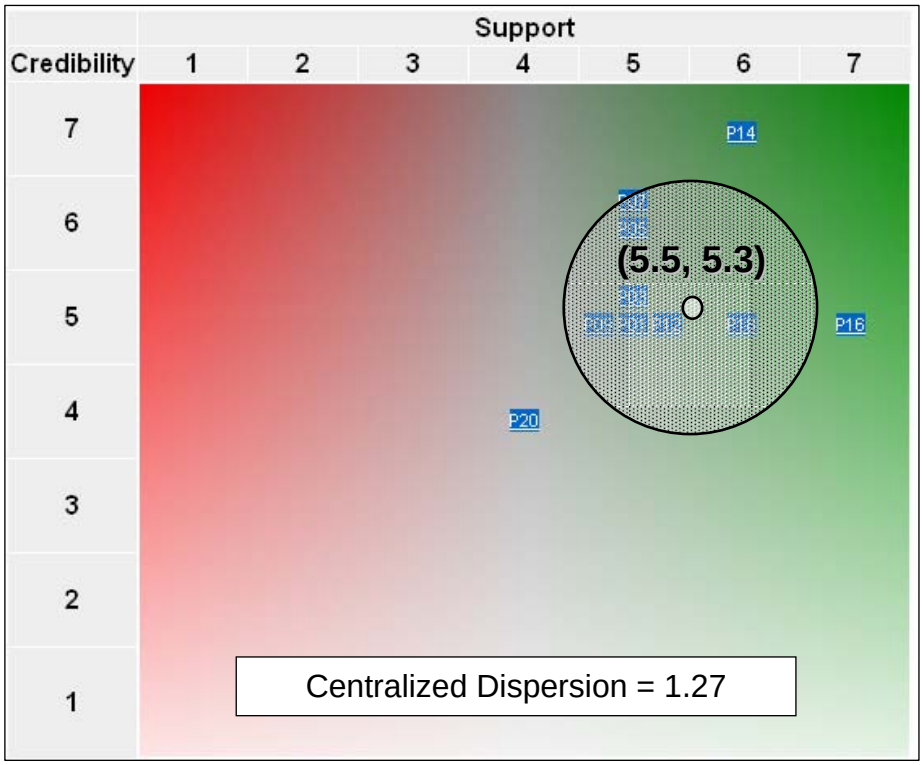

Figure 7: Centralized discussion interface with a dispersion of 1.27

Experimental Hypotheses

1. Dispersion would decrease over time with the decentralized as compared to the centralized discussion architecture because the discussion forum would facilitate reaching consensus among participants over time.
2. There would be greater "clumping" and global realignment of I-Post positions for those using the decentralized as compared to the centralized architecture because of more opportunity for social comparison using an open discussion forum.
3. There would be a greater change in absolute I-Post position and increased frequency of movement in the presence as compared to the absence I-Scape arrangement because the ability to view others' I-Posts. This also served as a control for biases associated with I-Post placement.
4. The frequency of I-Post movement would be positively correlated with discussion frequency as a result of information exchange in the discussion forum resulting in iterative reassessments of I-Post positions. However, results demonstrated that participants completed the near-minimum requirements for discussion and I-Post assessments; thus, discussion frequency and I-Post assessment frequency were similar across conditions.

Results

Time

It was expected that there would be a change over time, as a result of collaboration, from initial to final I-Post assessment. As measured by I-Post position, the degree of support significantly increased over time ($\eta p^2 = .34$) where Initial I-Post placement was 4.4 as compared to Final I-Post placement of 4.8. That is, I-Posts supported the hypotheses at their initial assessment and were rated as even more supportive at the final assessment. However, the credibility of the source did not change over time. Interestingly, the results also revealed that the mean position of the I-Posts was greater than the midpoint of 4 on the scale for both support and credibility ratings at both initial and final assessments. Collectively, the net positivity in these I-Post position results is suggestive of the confirmation bias at work.

I-Scape Arrangement

Unexpectedly, there was no main effect of I-Scape arrangement on I-Post position. That is, the presence or absence of others' I-Posts had no effect on I-Post placement. The mean I-Post position on the support scale was 4.6 in the presence and 4.7 in the absence condition. For the credibility scale, the mean I-Post position was 5.3 in the presence and 5.2 in the absence condition. There was also no effect of I-Scape arrangement on I-Post position over time. The lack of any arrangement effect may be due to participants completing the near-minimum posting requirements. Further, there may have been a ceiling effect from the confirmation bias which restricted the range of I-Post movement on the already positive I-Scapes.

Time x Architecture

However, there was an interaction between time and discussion architecture as measured by dispersion. As shown in Figure 8, dispersion *decreased* within *decentralized* and *increased* in *centralized* discussion architectures over time ($\eta p^2 = .24$). As expected, *decentralized* discussions began to consolidate I-Post positions over time – a likely result of the computer-mediated social comparison and discussion content.

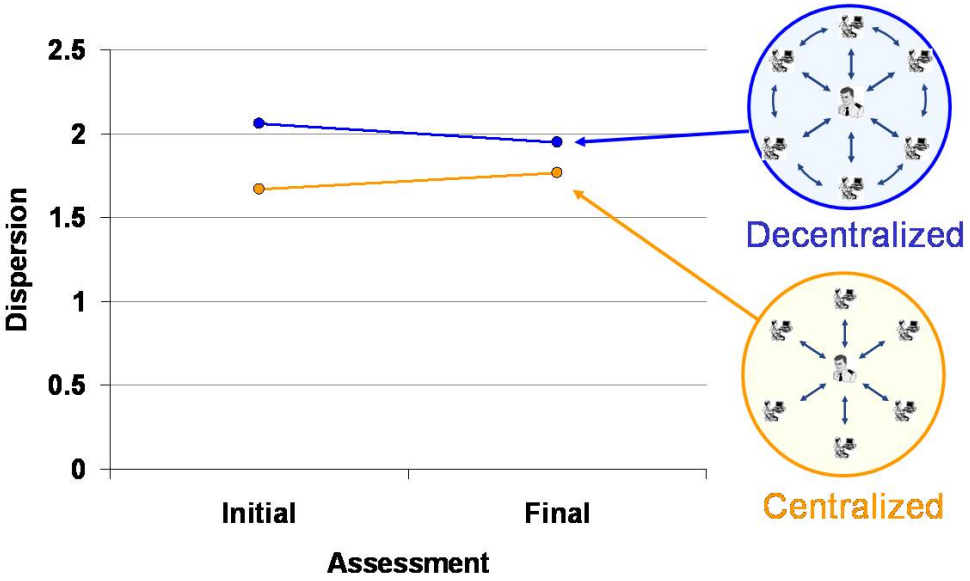

Figure 8: A Time x Architecture interaction
for assessment consensus over time.

Hypothesis Response Agreement

As previously mentioned, the final phase of the study required participants to write a short summary stating their position on their own hypothesis. Recall that the same hypotheses were used in both discussion architectures and hypothesis owners could always see others' I-Posts (presence). When hypothesis responses were compared between discussion architectures, 55% of hypotheses were in agreement (supporting or refuting), thus 45% of hypotheses were in disagreement between discussion architectures (e.g. centralized hypothesis owner supported hypothesis #1 while decentralized hypothesis owner refuted hypothesis #1).

Furthermore, 75% of those hypotheses in disagreement, the *centralized* discussion architecture supported whereas the *decentralized* refuted the hypothesis. Of note, 67% of hypotheses were *refuted* by at least one owner even though there was a confirmation bias for both support and credibility scores. In other words, the majority of hypothesis owners refuted the hypothesis even when the mean ratings of others' I-Posts were in support of the hypothesis. Collectively, these results suggest evidence of confirmation bias, where people selectively seek information to confirm hypotheses. Therefore, this may also

suggest that the discussion content may have had more impact on final decision-making than initially anticipated.

Conclusion

JIGSAW was used in a unique web-based collaboration experiment to solve intelligence problems *in vivo* over four weeks. The study also demonstrated how applied research was able to inform design and integrate into the prototype design cycle. That is, future instantiations of JIGSAW will be able to leverage the results and tailor discussion architecture appropriately into the system.

A major contribution of this study was the novel metrics for collaborative assessment using a measure of dispersion that is capable of assessing the dynamic process of group collaboration. Regardless of discussion architecture, collaborating analysts were still prone to the confirmation bias. The results of the present study are consistent with results obtained by Cook and Smallman (2007) at the individual level. Conceptually, it appears these biases originate from the micro/individual level and persist up to the macro/team level as represented in Figure 10.

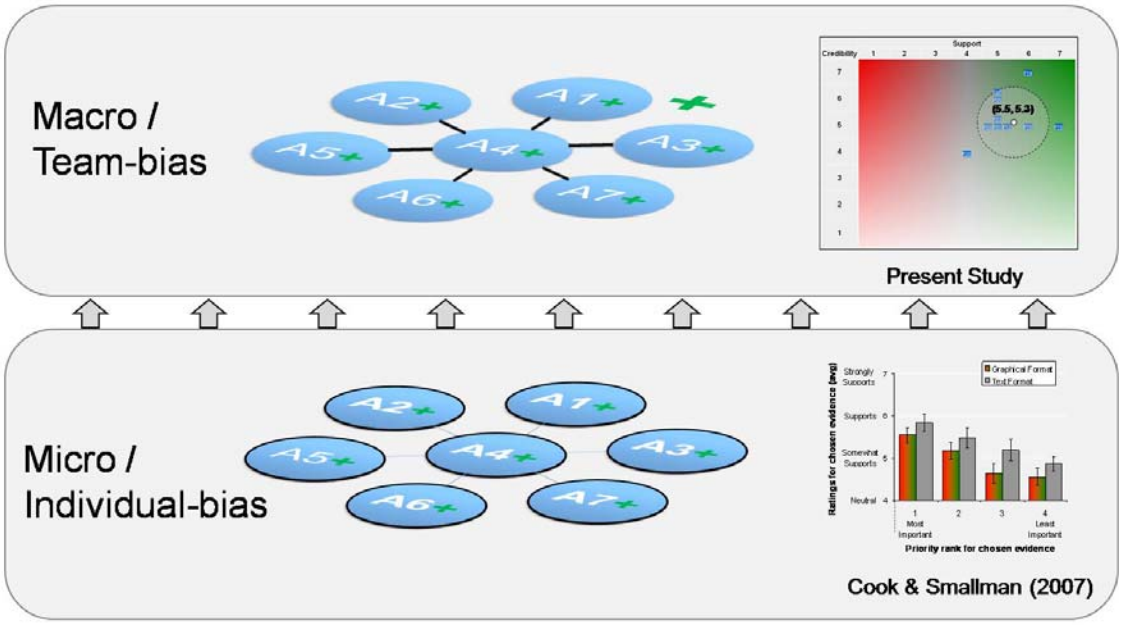

Figure 10. Representation of the relationship of the pervasive confirmation biases observed at both the individual (bottom) and team (top) levels

However, discussion architectures did function differently over time as demonstrated by the measure of dispersion. That is, participants role-playing collaborating analysts reached consensus faster with decentralized discussion. Decentralized methods of collaboration facilitated the exchange of information towards a unified viewpoint. The mechanism of action for this effect is likely due to the computer-mediated social

comparison that takes place via a shared visual workspace and discussion content. The discussion and information content had a greater impact on assessment than display arrangement. Therefore, it may be necessary to temporarily mask information content from new JIGSAW users so as not to bias their novel contribution or interpretation of the I-Scape and discussion forum. However, collaboration methods and/or visualization techniques of the I-Scape interface still need to address the pervasive confirmation bias.

As mentioned above, the dispersion metric is a scalar quantity that does not indicate the direction of change over time (e.g. analysts can not readily observe the shift towards positivity). Not only is the dispersion metric useful as a way to characterize collaborative assessment behavior, it could have applicability as a decision support concept in itself. For example, military analysts and collection managers have expressed to us a desire for the state of assessment on an I-Scape to be summarized so that temporal trends are more salient. In the present study, the dispersion metric was only used for analysis and was not graphically presented to participants. For example, Figure 11 shows one promising idea of a "compass rose" visualization to turn the dispersion metric into a vector quantity and support situation awareness of analysis trends. Using the concept of the dispersion metric, it shows the center of gravity of the I-Posts on the I-Scape (the center of the rose), the dispersion of the I-Posts (the circular extent of the rose), but also includes an arrow to indicate the prevailing trend over some user-indicated time period (the arrow). Such a visualization would provide at-a-glance situation awareness of trends – an awareness that is almost impossible to quickly glean using current analyst tools and procedures. The rose is shown with some user-defined drop-down options depending on their interests and requirements.

Figure 11: Harnessing the dispersion metric for use in an I-Scape temporal trend visualization.

Collectively, these data suggest that intelligence analysts would benefit from collaborative tools that support information-sharing from unknown sources in a common workspace. By taking a macro-cognitive perspective that emphasizes team-level collaboration, we have been able to suggest novel, simpler interventions and metrics than those that have emerged from previous and more micro-cognitive analyses. JIGSAW represents the application, integration, and now transition of collaborative science and technology concepts, many of which were researched and developed in the Collaborative Knowledge Interoperability (CKI) research program of the Office of Naval Research (ONR). JIGSAW has the potential to begin to transform the way analysis is performed, making it more systematic, responsive, and collaborative.

References

Cook, M.B., & Smallman, H.S. (2007). Visual evidence landscapes: Reducing bias in collaborative intelligence analysis. In *Proceedings of the 51st Annual Meeting of the Human Factors and Ergonomics Society*, (pps. 303 – 307), Human Factors and Ergonomics Society, Santa Monica, CA.

Cook, M.B., & Smallman, H.S. (2008). Debiasing analysis using visual evidence landscapes. *Human Factors*, submitted.

Hall, T. H. (1998). Collaborative computing technologies and hierarchical command and control structures. Paper presented to the *3rd International Command and Control Research and Technology Symposium*. Naval Postgraduate School, Monterey, CA, June 29th- July 1st.

Heacox, N.J., Moore, R. A., Morrison, J. G., & Yturralde, R. F. (2004). Real-time online communications: 'Chat' use in Navy operations. In Proceedings of the 2004 *Command and Control Research and Technology Symposium.* San Diego, CA: Space and Naval Warfare Systems Center, 15-17 June 2004.

Heuer, R. J. Jr. (1999). *The Psychology of Intelligence Analysis*. Center for the Study of Intelligence, Washington, DC.

Hogarth, R.M. (1987). *Judgment and choice: the psychology of decision making*. 2nd Edition. New York, NY: Wiley.

Hutchins, S. G., Pirolli, P. L., & Card, S. K. (2004). A new perspective on use of the critical decision method with intelligence analysts. In *Proceedings of the 2004 Command and Control Research and Technology Symposium.* San Diego, CA, 15-17 June 2004.

Janis, I. L. (1982). *Groupthink* (2nd revised Ed.). Boston: Houghton Mifflin.

Letsky, M.P., Warner, N., Fiore, S.M., Rosen, M., & Salsa, E. (2007). Macrocognition in complex team problem solving. In *Proceedings of the 2007 International Command and Control Research and Technology Symposium.* San Diego, CA, 15-17 June 2007.

Nickerson, R.S. (1998). Confirmation bias: a ubiquitous phenomenon in many guises. *Review of General Psychology*, **2**, 175-220.

Patterson, E.S., Woods, D.D., Tinapple, D., & Roth, E.M. (2001). Using cognitive task analysis (CTA) to seed design concepts for intelligence analysts under data overload. In *Proceedings of the 45th Annual Meeting of the Human Factors and Ergonomics Society*, Minneapolis, MN: October 8-12, 2001.

Smallman, H.S. (2008). JIGSAW – Joint Intelligence Graphical Situation Awareness Web for collaborative intelligence analysis. In *Macrocognition in Teams*, Letsky, M.P., Warner, N., Fiore, S., & Smith, C.A.P. (Eds). Ashgate Publishing, in press.

St. John, M., Gwynne, J.W., Kelly, R.T., & Smillie, R.J. (2005). ISR Reach-Back: A human-systems integration assessment from Trident Warrior 2004. Paper presented at the 10th *International Command and Control Research and Technology Symposium*, McLean, VA, June 13-16, 2005.

The 9/11 Commission (2004). *The 9/11 Commission Report: Final Report of the National Commission on Terrorist Attacks Upon the United States*. Washington, DC.

Warner, N. W., Letsky, M. P., & Cowen, M. B. (2005). Cognitive model of team collaboration: macro-cognitive focus. In *Proceedings of the 49th Annual Meeting of the Human Factors and Ergonomics Society*, (pp.269-273). Human Factors and Ergonomics Society, Santa Monica, CA.