CONVERGING APPROACHES TO AUTOMATED COMMUNICATIONS-BASED ASSESSMENT OF TEAM SITUATION AWARENESS

Shawn A. Weil, Pacey Foster, Jared Freeman
Aptima, Inc.

Kathleen Carley, Jana Diesner, & Terrill Franz
Carnegie Mellon University

Nancy J. Cooke, Steve Shope & Jamie C. Gorman
Cognitive Engineering Research Institute

Introduction

Collaboration enables people to execute tasks that are beyond the capabilities of any one of them. Each member of an organization or team has a set of skills, roles, and responsibilities that, when executed accurately and in synchronization with other members of that organization or team, enables them to accomplish the work of the organization. In complex, dynamic environments, team collaboration is more than the simple aggregation of the work products of individuals; collaboration also requires complex exchanges of information, largely through spoken or written language. Linguistic communication is the choreography of team performance.

Modern networked information systems support synchronous and asynchronous communication among globally distributed team members via telephone, e-mail, instant messaging, and text chat-rooms, making possible the coordination of activities that would have been impossible or impractical in the past. The distributed project coordination common in commercial organizations would not be tenable without these technologies.

However, these technologies do not ensure highly effective organization. Cognitive collaboration quality varies significantly between teams, whether they are collocated (McComb, 2005; Isaacs and Clark, 1987; Warner & Wroblewski, 2004), or distributed over networks or in virtual environments (Cooke, 2005; Entin and Serfaty, 1999; Cooke, Gorman, & Kiekel, this volume). When collaboration tools do make a difference, it is not always positive. These tools increase the opportunity for information overload; errors of commission (miscommunication) can increase relative to errors of omission (non-communication); and decisions and actions sometimes – and sometimes tragically – are mis-coordinated (Woods, Patterson, & Roth, 2002, Weil et al, 2004).

The Department of Defense (DoD) uses networked collaboration technologies to coordinate distributed, heterogeneous forces for both wartime and peacetime activities. Such coordination is a key element of Network Centric Warfare or Network Centric Operations (NCO), a theory of warfare in the information age. Among the primary tenets of NCO is the belief that networked information and collaboration increases “shared situation awareness”, a common understanding of the state of the mission environment. This, in turn, is predicted to enhance the effectiveness of forces (Alberts, 2002). Some critics foresee a concurrent increase in the errors cited above. The debate is important to the nation’s economic health, as well as its defense. However, NCO concepts are not the sole province of military organizations; they are also being applied in large and economically important commercial organizations. Effective work product requires that the individual contributors align their understanding of the state of the commercial environment, the corporate mission within it, and their responsibilities. Tools such as data dashboards, knowledge portals, distributed conferencing applications, and chat are often installed to facilitate this. Corporations, like DoD, are putting NCO concepts and technology to use to build situation awareness.

DoD has increased funding to develop common operating picture displays, collaborative environments, and communication tools, explicitly to improve situation awareness. However, aspects of situation awareness are not often considered when these tools are developed. Nor have there been many attempt to assess whether, how, or how much these tools improve situation awareness. Thus, there is little evidence from which to determine whether and how to (re)design the tools, the organizations that use them, or the procedures for doing so.

To address this gap in knowledge, the authors are developing IMAGES – the Instrument of the Measurement and Advancement of Group Environmental Situation awareness. IMAGES is a software tool that gives mission personnel and researchers access to multiple, complementary communications assessment techniques that provide insight into both the content and process of collaboration in an organization. The capability to analyze content should enable leaders of organizations and designers of systems to measure shared situation awareness unobtrusively (Weil, 2006). The capability to assess collaborative process should enable them to manage the development of SA. In this chapter, we describe the approaches used to assess team performance via analysis of communications rather than describing the software.

Situation awareness and Macrocognition

Individual, Team, and Shared Situation awareness

Effective collaboration and the utility of the NCO doctrine are predicated on the concept of situation awareness. Endsley (1988) defines three levels of individual *situation awareness*: (1) perception of elements in the environment, (2) comprehension of meaning in those elements, and (3) use of that understanding to project future states. High situation awareness is associated with good performance, as the individual can anticipate the actions of elements in his/her environment. Low situation awareness is not desirable, as actions taken by the individual will likely be inappropriate given the true state of the environment.

When assessing an organization rather than an individual, situation awareness requires further definition. Organizational situation awareness is not simply the aggregate situation awareness of all constituent members. Endsley has distinguished between *shared situation awareness* (shared SA) – "the degree to which team members possess a shared understanding of the situation with regard to their shared SA requirements (Endsley & Jones, 1997)" – and *team situation awareness* (team SA) – "the degree to which every team member possesses the SA required for his or her responsibilities (Endsley, 1995)." In complex situations, it is often counterproductive for every member of an organization to have total knowledge of the state of the environment (i.e., completely shared SA) – nor is it advantageous for information and knowledge to be perfectly partitioned among members of the organization without overlap (i.e., perfect team SA). For teams, a balance of shared and team SA is required (e.g., Cooke, Salas, & Cannon-Bowers, 2000). Discovering the right balance for a given team – and measuring SA to determine whether that balance is being met – is an unsolved problem. It is not simply a matter of "more is better."

It is equally challenging to identify the communication processes that produce a given level of SA balance of SA types (shared vs. team). This diagnosis of process is required to prescribe a remedy, a change to communication processes that improves SA in the long run.

Individual situation awareness is a largely internalized construct. The outward behaviors of individuals may reflect some aspects of situation awareness, but a true empirical evaluation typically requires interviews, surveys, or highly instrumented and artificial experiments. However, the elements of *collaboration* that produce team and shared SA are external; they are evident in communications content and process (Cooke, 2005). Thus, they are observable and measurable in operational settings, without the intrusive methods cited above.

Macrocognition

The approach to assessing situation awareness described in this chapter relates to several of the macrocognitive processes described elsewhere in this volume (Warner, Letsky, Wroblewski, 2007)The behaviors of teams and organizations—both externalized behaviors associated with Joint Cognitive Systems and internalized events related to the perception and cognition of individuals—have been associated with the construct of Macrocognition (Warner et al., 2005; Klein et al., 2003; Cacciabue & Hollnagel 1995). Macrocognition describes the cognitive functions that are performed in natural decision-making settings, typically by multiple, interacting individuals. Warner et al., 2005 have attempted to formalize this construct, positing four collaboration stages—Knowledge Construction, Collaborative Team Problem Solving, Team Consensus, and Outcome Evaluation and Revision. Within these stages are between 10 and 20 distinguishable macrocognitive processes that relate to the process of problem solving, mental model creation, and information exchange.

The externalized communication behaviors posited to be related to organizational SA relate to several of the macrocognitive processes described in Warner et al (2005). The process of consensus building (the macrocognitive process of *knowledge interoperability development*) is supported by linguistic interchange among members of a collaborating team – the same communications content we will use to assess SA. The *team shared understanding development* macrocognitive process reflects some of the distinctions between shared and team SA described in the preceding section. The implied, cyclical form of the model is consistent with the notion that communications process (not just static content) predicts overall trends in SA.

However, the level of aggregation implicit in the large organizations on which we are focused – dozens to hundreds of individuals – seems to be incompatible with some of the macrocognitive processes. For example, the macrocognitive process "*convergence of individual mental models to team mental model"* posits a single mental model for the group. We argue that this is neither possible nor advantageous in large groups. Instead, we advocate for individual mental models or situational assessments that are aligned with the tasks required.

Whether developing a theory of macrocognition or designing technology to facilitate macrocognition, assessment is central. To illustrate the point, consider the concept of human intelligence. Without assessment in the form of various intelligence quotient (IQ) tests, we would have very little to say about a theory of intelligence and no evidence as to whether interventions improve intelligence or not. Likewise, assessment of macrocognition is necessary for theory development and for testing the success of interventions designed to facilitate it.

Two forms of assessment are required to understand and manage SA. Assessments of the state of SA are descriptive; they are valuable for comparing and evaluating new technologies, techniques, and organizational forms prior to acquisition or implementation. Assessments of the process of SA are diagnostic; they enable us to understand the root causes of success and failure of macrocognitive processes (communication processes, in particular) and to manage those processes. The methods and tools described in this report focus on communication assessment as a means of assessing macrocognition.

Multiple communications assessment approaches

In the research literature, communication analysis methods focus either on communication content (i.e., what is said; Carley, 1997; Kiekel, Cooke, Foltz, Gorman & Martin, 2002; Kiekel, Cooke, Foltz, & Shope, 2001), communication flow (who talks to whom; Kiekel, Gorman, & Cooke, 2004), or communication manner (style, intonation, gestures, face-to-face or not, etc.; Queck, et al., 2002). Communication content and communication flow have been the primary focus of team researchers.

Content methods attempt to summarize the content of discourse and then associate the content with team performance or macrocogntive processes. Although there have been successes in applying content-based communication analysis techniques to the understanding of teams, there are some drawbacks including the need to translate the spoken record into words (i.e., transcription) either manually (a tedious and resource-intense process) or by speech recognition software (an inexact process). In some cases the terminology used is domain-specific such that an ontology for that domain must be developed prior to interpreting the transcribed terms. Once the transcription and initial domain analysis is complete, the content-based methods pay off by providing a good source of information about the content (the state) of macrocognition.

Another source of information that is perhaps more tied to macrocognitive processes is the analysis of communication flow. Specific interaction or flow patterns in communications have been associated with effective team performance or team SA (Kiekel, Gorman, & Cooke, 2004). Similarly some patterns may signal a loss of team situation awareness (Gorman, Cooke, Pedersen, et al., 2005; Gorman, Cooke, & Winner, 2006). Algorithms and software have been developed to extract patterns in the specific timing and sequences of interactions. Although the data used in this analysis is very basic, it is easy and inexpensive to collect and analyze and is conducive to automation. Thus there is a tradeoff between ease of analysis and potential automation, and richness of the resulting data.

The authors have taken a multi-faceted approach to measurement of situation awareness through the assessment of communications. The first employs a set of networks aimed at representing both content and interactivity in an organization. The second is inherently temporal in nature, relying on regularities in turntaking behavior in communication. These two methods are described in detail below.

Network Representations of Content and Structure

As speakers of a language, humans can often read a transcript or listen to a recording of an event to understand the meaning of interchanges, follow conversational threads as the propagate through an organization, surmise the level of situation awareness within the team communicating, and interpret the relationships between issues being discussed and action in the world. Distilling this content into forms that lend themselves to measurement and analytic interpretation would allow greater insight into situation awareness in large, complex organizations.

Several representations of the content and structures that underlie an organization's behavior can be extracted from standard communications media. Transformed into matrices, metrics can be developed that can aid in interpretation. In the current effort, we use the approach developed by Carley (2002) in which entity extraction techniques, embedded organizational ontologies, and semi-automated thesaurus construction are used to extract and link both social networks and semantic networks. In this *meta-matrix* *approach,* the social structure of an organization is derived from mention of pertinent organizational elements (e.g., agents, resources, locations), within a corpus of texts. These components are classified into an ontology that structures them into a model of a social system. This model allows investigation of the composition of a social system, and identifies the connections among organizational components. In a given domain, this type of inquiry becomes highly automatized; models of organizations can quickly be extracted from large amounts of text, whereas manual extraction of an organizational model from text would be laborious and error prone. The meta-matrix scheme, derived from organizational research, provides such ontology for modeling the social and organizational structure of teams, groups, distributed teams, clans, organizations and so on (Carley, 2003, 2002; Krackhardt and Carley, 1998). In the meta-matrix approach the entities of a social system are agents, organizations, knowledge, resources, tasks and events and locations. Previously, Diesner and Carley (2005) have described an approach for combining networks that reflect the content of communications (i.e., map analysis) with the meta-matrix model. The resulting integrative technique we refer to as Meta Matrix Text Analysis (Diesner & Carley, 2005). This technique enables analysts to extract not only knowledge networks, but also social and organizational structure from texts.

We have identified five of network types that may provide insight into organizational situation awareness and other macrocognitive processes.

I. The Knowledge Network The *knowledge network* is a representation of the externalized knowledge or understanding of those individuals communicating. One way to represent the semantic content of communication is through the use of Network Text Analysis (NTA; Popping, 2000). NTA is based on the notion that language and knowledge can be represented as a network of words and the relations among them (Sowa, 1984). Map analysis (Carley, 1993, 1997b; Carley and Palmquist, 1992), as achieved using AutoMap (Carley, Deisner & Doreno, 2006) is one method to create these semantic networks. The major concepts in a text are extracted and become nodes in a network. These nodes are construed as "concepts." The arcs among the nodes – or the relationship among concepts – are defined by the proximity of those concepts to each other in the text. Pairs of concepts are construed as "statements." Given sufficient text, a complex web or network of concepts and statements is created (Carley, 1997b).

This *knowledge network* reflect some of the complexity of the semantic and syntactic structure of the original texts, but in a form that allows for easier manipulation and perhaps automated interpretation. One way to construe these *knowledge networks* are as proxies for the aggregate knowledge of the individuals who framed the communication. As communication is the observable engine of team cognition, the semantic networks based on that communication become a representation of team cognition and/or organizational understanding. By selecting and comparing the semantic networks for different groups of people within and organization and specified time periods, you can assess the similarity or divergence of the knowledge of the individuals communicating. It is in this way that team and shared situation awareness can be automatically derived. In the parlance of the meta-matrix, this is known as a "knowledge x knowledge" network, as all of the nodes are given the "knowledge" attribute.

II. The Social Network A *social network* describes the relationships among individuals in an organization. In a corpus of e-mails, a social network can be easily generated from the "To," "From," and "CC" lines. This social network is representation of interaction among members of an organization. Within the meta-matrix ontology, each of the nodes is given the attribute "agent," and the resulting network becomes an "agent x agent" network.

III. The "Agent x Knowledge" Network Combining the *social network* and the *knowledge network* allows inquiry into the relationship between patterns of interaction and the resulting change in situational understanding. As the meta-matrix model allows nodes to be identified by ontological type or class, a new "*agent x knowledge" network* can be created in which some nodes represent the individuals communicating (i.e., "agents"), while other nodes represent the "knowledge" being communicated by those individuals.

IV. The Implicit Meta-Matrix While the text of communication can be distilled into a knowledge network, many of the nodes in that network can be further described using the meta-matrix attribute labels (e.g., location, agent, task, event, etc). A new network representation is thus created, an *implicit meta-matrix*, which adds additional specificity and affords additional measurement possibilities. We have used the term *implicit* because the relationships among the nodes does necessarily refer to real world connections among the nodes (as is the case in the *social network* described above) but instead could be construed as the aggregate mental model of those communicating.

V. The "Agent x 'implicit meta-matrix' network." Finally, a network can be created in which the *social network* derived from message headers is combined with the *implicit meta-matrix* derived from the content of messages. The resulting network is an "*agent x 'implicit meta-matrix' network.*" As in "*agent x knowledge" network* above (III), the combination of a network based on patterns of interactivity with one representing content allows researchers and operators to gauge how knowledge is affected by different types of interactivity. This "*agent x 'implicit meta-matrix' network*" allows more nuanced assessment of changes in knowledge, as both the change in meta-matrix nodes and the structural arrangements of those nodes can be assessed.

Use of Meta-matrix Networks The purpose for creating these networks is to enable inquiries into the state of knowledge of an organization at a given time and the relationship of those states to patterns of interactivity among members of that organization. The degree of shared situation awareness between two groups within an organization can be determined by measuring the degree of overlap in the Knowledge Networks or Implicit Meta-Matrices created for each of those groups. Comparing networks derived from the same individuals at different time periods provides insight into organizational change, and can be used in conjunction with knowledge of the world to correlate patterns in communication/interactivity with real-world consequences.

Flow Analysis

The network approaches described in the previous section emphasize content and interactivity. However, they are relatively static in nature – dynamism is implied only when comparing networks based on different time spans. However, communication flow is inherently dynamic. To complement the network approaches, we have also included an explicitly temporal approach to assessing situation awareness. One set of algorithms that has been developed to process flow data is called FAUCET (Flow Analysis of Utterance Communication Events in Teams; Kiekel, 2004). FAUCET metrics have been developed and validated in a UAV command-and-control scenario and more recently in the context of Enron emails. One of the FAUCET metrics is called ProNet (Cooke, Neville, & Rowe, 1996). ProNet is a method for reducing sequential event data that relies on the Pathfinder algorithm (Schvaneveldt, 1990) as its kernel. ProNet has been recently applied to communication data with some success (Kiekel,2004). The result of this analysis is a graph structure in which communication events that occur together frequently are connected by directed links. ProNet, like Pathfinder, is limited in the sense that the multiple-link paths represented in the network structure are only certain to exist on a pair-by-pair basis. ChainMaster is a software tool that implements the ProNet algorithm, but that extends it by doing tests for the existence of chains at multiple lags. With these tests, the likely multiple-link paths can be highlighted. More information about these approaches can be found in Cooke, Gorman, and Kiekel (2007).

Assessing Situation awareness in Large Organizations: The Enron Example

The purpose of the current effort was to explore ways of automatically assessing situation awareness of large organizations using converging, complementary communications assessment methods. To illustrate this assessment, we required a suitable corpus of communications. As most military corpora are classified, we chose a corpus of e-mail from a publicly available corporate entity, the Enron corporation. This corpus was a reasonable approximate of the types of corpora we would expect to see in large military organizations: there are several hundred people interacting, there are hundreds of thousands of messages over a multi-month period, and there are observable events in the public record to correlate with patterns in the data. This section describes the Enron corpus, our assessment methodology, the preliminary results, and a high level interpretation

The Enron Accounting Scandal

The Enron accounting scandal of 2001 was one of the largest (and most widely known) cases of corporate malfeasance in U.S. history. Beginning in the summer of 2001, revelations about the scope and extent of Enron's accounting practices started to become public. While the events precipitating the collapse had started many years before with a series of illegal accounting practices, the fall of 2001 represented the beginning of the end for Enron. Between August and December 2001 Enron CEO Jeffrey Skilling resigned, a series of Wall St. Journal articles reported on the the scandal, stock prices fell from over $80 per share in January 2001 to junk bond status by January 2002, and the Security and Exchange Commission launched a formal investigation into Enron's accounting practices.

These calamitous events culminated in Enron's declaration of bankruptcy in December, 2001. In the end, the Enron accounting scandal wiped out $68bn of market value and caused irreparable damage to investor confidence as well as eliminating over $1bn of employee retirement funds held in Enron stock. The aftermath of this crisis sent shock waves through the stock market and has led to sweeping legal and regulatory changes such as the Sarbanes-Oxley act as well as many years of litigation.

An important product of the Enron investigation was a large corpus of e-mail communication among Enron executives and employees prior to its bankruptcy in December 2001. Because it represents a rare look at the real-time communication inside an organization as it managed a life threatening crisis, the Enron e-mail corpus is ideally suited to the development of automatic tools for real time assessment of group level situation awareness.

The Enron e-mail corpus was originally released by the Federal Energy Regulatory Commission (FERC) during its independent investigation. It captures data extracted from the e-mail folders of 151 Enron employees obtained during the FERC investigations. However, because of the ad hoc nature of the sample and the raw format of the data, the raw corpus presented challenges for subsequent researchers. Some of these challenges were related to the unstructured nature of the data. For example, employees with multiple e-mail addresses appear in the original data as different people while some e-mail addresses (such as automatically generated responses from servers) do not represent people at all. Several different research groups have addressed these problems in various ways leading to a proliferation of versions of the Enron corpus over time.

Deisner, Frantz & Carley (2005) provide a comprehensive overview of the iterations of this data set which we will not duplicate here. Suffice it to say that the primary differences between Enron corpora have to do with the methods used to clean the data. For the purposes of the current study, we obtained a later generation of the Enron corpus which was cleaned by researchers at the University of Southern California and placed in a SQL server. Each of the research teams used a sub set of this corpus for their work providing triangulation among multiple views of the situation awareness of Enron employees during the final months of the company's life.

The Enron Corpus

The Federal Energy Regulatory Commission (FERC) gathered 619,449 emails from the computers of Enron employees, mainly senior managers[1], as part of their investigation. In May of 2002, FERC publicly released the Enron email dataset. For each email, the email address of the sender and receiver(s) were made available, as well as the email's date, time, subject and body, while attachments were not released. The emails were written between 1997 to 2004.

---

[1] Note that the Enron email corpus contains a plethora of emails written by individuals who were not involved in any of the actions that were investigated by FERC.

SRI International (SRI), Cohen from CMU, Corrada-Emmanuel from the UMASS, Shetty and Adibi from the ISI group at USC, UC Berkeley Enron Email Analysis Project, and Corman and Contractor together with further members of the Organizational Communication Division of the International Communication (ICA) each played a hand in reducing and refining the Enron dataset and providing interfaces to the data. Their resulting database consists of 252,759 emails in 3000 user defined folders from 151 people.

Methods

As described above, research on situation awareness has been largely confined to studies on small groups in command and control settings. Our research seeks to broaden the construct of situation awareness to measure the impact of critical events on knowledge and information sharing at the organizational level. Because the measurement and assessment of situation awareness at the organizational level represents new theoretical and methodological terrain, we used inductive, exploratory methods to understand how critical organizational events affected the content and structure of communication during the fall of Enron.

Our basic research question reflects one of the central theoretical assumptions of small group research on situation awareness, namely, that critical events will serve as orienting stimuli that generate increased communication and information sharing as collaborating groups seek to make sense of and coordinate reactions to unexpected events. While this has been a focus of many small group studies on situation awareness, it has not been addressed at the organizational level. To explore this notion at the organization level, we selected a time period containing five critical events in the decline of Enron and compared the structure and content of organizational communication around these events with the communication in times that did not contain critical events. Table 16.1 below depicts five critical events that took place at Enron between August and December of 2001 including the resignation of Jeff Skilling, the announcement of the SEC investigation and the declaration of bankruptcy in December. Each of the five events we selected was large, widely known by members of the organization, and likely served as a focal point of attention and corporate communication. To understand how organizational situation awareness changes around critical events, we also selected five dates between August and December 2001 when there were no critical events at Enron[2]. Using the three business days before and after these critical and non-critical dates, we are able to compare how the structure and content of organization-wide communication changes in response to critical events.

Table 16.1 Critical events at Enron from August – December 2001

Date	**Critical Events (Non Events) at Enron**
August 1, 2001	(No event)
August 14, 2001	Jeff Skilling announces his resignation
September 11, 2001	(No event)

[2] Readers might note that September 11, 2001 is one of these "non critical" dates. We selected this date because it represents a useful quasi-control as it was clearly a critical event in the history of the world, but was not directly related to Enron per se. This allows us to explore whether organization level situation awareness changes in response to critical internal organizational events as opposed to external events.

September 25, 2001	Ken Lay assures employees that Enron's accounting practices are legal
October 2, 2001	(No event)
October 16, 2001	Enron announces SEC investigation
October 23, 2001	(No event)
November 6, 2001	Enron announces profit overstatement
November 20, 2001	(No event)
December 4, 2001	Enron files for bankruptcy

Results

Semantic and Social Network Analyses of the Enron e-mail Corpus

This section describes the results of the network analysis of the Enron e-mail corpus. We begin by describing the data pre-processing that was performed followed with a description of the measures that were calculated to measure changes in group level situation awareness at Enron.

Before texts are analyzed, they can be pre-processed in order to normalize the data and to reduce the data to the terms relevant for a research question. For this particular project a series of pre-processing steps were required to make the e-mail corpus suitable for network analysis. All the texts were automatically cleaned to remove non-content bearing symbols, such as apostrophes and brackets. Text that was to be used in network text analysis (e.g., all the content of e-mails, but not the header information) was submitted to several additional preprocessing steps.

First, we performed deletion, which removes non-content bearing conjunctions and articles from texts (Carley, 1993). The delete list we built and applied list was tailored for this data set and contained 32 entries. Second, we built and applied thesauri are used to resolve synonyms and abbreviations. When applying a thesaurus, AutoMap searches the text set for the text-level words denoted in the thesaurus and translates matches into the corresponding words. Because the terminology of a thesaurus depends on the content and the subject of the data set we used a thesaurus developed specifically for analysis of the Enron e-mail corpus.

Our third step was to construct and apply a generalization thesaurus. A generalization thesaurus typically is a two-columned collection that associates text-level concepts with higher-level concepts, e.g. in order to convert multiple noun phrases into single noun phrases (such as Jeff Skilling into Jeff_Skilling). The text-level concepts represent the content of a data set, and the higher-level concepts represent the text-level concepts in a generalized way (Popping & Roberts, 1997). The generalization thesaurus we built contained 517 association pairs. To ease the construction of the generalization thesaurus, we also performed Named-Entity Recognition in order to automatically identify names of people, places and organizations, as well as bi-gram detection, which return the most frequent combinations of any two terms in a data set. After this stage of pre-processing, semantic network analyses were run in order to extract Knowledge networks (KK).

Our fourth and final pre-processing step was to construct a Meta-Matrix Thesaurus. This thesaurus is needed in order to perform meta-matrix text analysis, which enables the classification and analysis of concepts in texts according to the Meta-Matrix model ontology (Diesner & Carley, 2005). Example: Jeff_Skilling will be associated with and translated into agent. Since one concept might need to be translated into several meta-matrix categories, a meta-matrix thesaurus can consist of more than two columns. The meta-matrix thesaurus we built associated 482 words with meta-matrix categories. After this stage of pre-processing, meta-matrix text analysis was performed in order to extract meta-matrix text networks from the data (iMM).

The CASOS Email Parser (CEMAP) enables the extraction of different types of network information from emails (e.g. who exchanges information, who provides what information, etc.). The following image shows what types of information can be extracted with CEMAP.

As described above, the social network (SN) represents social network data that can be extracted from email headers. This includes agent-agent networks, where agents are the people who sent and received an email, agent-task networks, where tasks are emails, and agent-knowledge networks, where knowledge is the content from the subject line. In SN, nodes represent people, and edges represent exchanged emails (frequency count). This network type does not require any text coding in AutoMap.

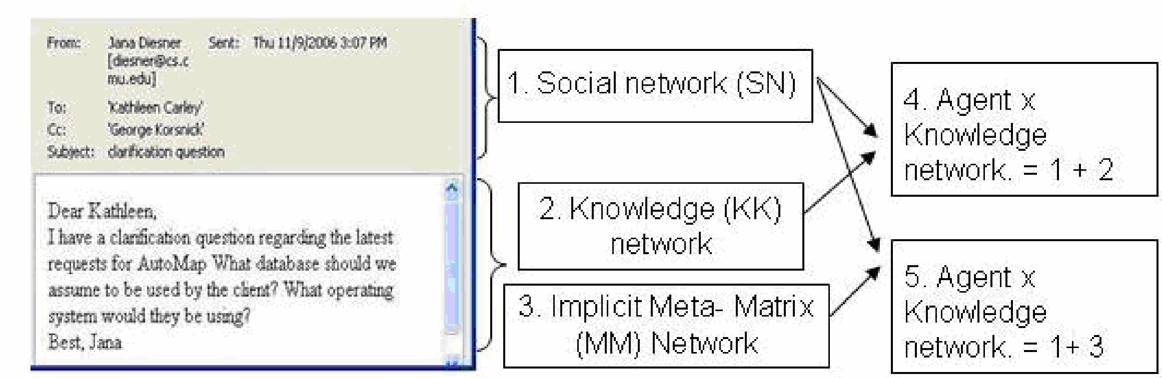

Figure 16.1 Network representations of e-mail data

The knowledge (KK) and implicit meta-matrices (iMM) are extracted by performing semantic text analysis with AutoMap (Diesner & Carley, 2004). More specifically, Knowledge networks (KK) represent semantic network or mental models that are contained in the bodies of individual emails. In KK, nodes represent knowledge items, and edges represent the co-occurrence of terms in text. For iMM, texts are coded in AutoMap according to a taxonomy or ontology (e.g. meta-matrix, while are ontologies can be specified by the user). In iMM, nodes represent instances of categories (e.g. agent, knowledge, resources) of the ontology, and edges represent co-occurrences of terms in texts.

Networks types 4 and 5 result from the combination of SN with KK and iMM, respectively. In type 4, nodes represent people and knowledge, and edges represent emails and mental models. In type 5, nodes represent the categories of the taxonomy as specified by the user, and the edges represent the co-occurrence of the terms that represent instances of the taxonomy in the corpus. For the creation of type 4 and 5, the extraction of type 2 and 3, respectively, is mandatory. CEMAP stores all network data as DyNetML files (Tsvetovat, Reminga & Carley, 2003) (a derivate of XML). This data can be analyzed with any package that reads DyNetML (e.g., ORA, Carley & DeReno, 2006).

To explore the impact of critical events on communication network structures at Enron, we generated multiple measures for each of the four matrices described above. A two-sample Welch's *t*-Test (variances assumed to be unequal) was conducted on each of the measures for the relevant, corresponding meta-matrices. Despite the use of multiple measures and multiple matrices, in no case could we reject the null hypothesis that the population means of the events and non-events samples were equal. For example, the communications network, which was derived from the email headers information, did not show a statistical difference in the clustering, number of components or cliques between the event and non-event samples. Similarly, the implicit organizational network which is derived from network text analysis of the email message content, does not show a statistical difference in any of the four change-measures inspected (clustering coefficient, weak component count, clique count the number of groups). The semantic network which is entirely based on the content of the email messages, does not show a statistical difference in any of the change-measures inspected. Finally, the results for both the inferred and semantic networks when each is combined with the communications network did not show a statistical difference in any of the four change-measures inspected. Likewise, The Communication and inferred Organization network, which is based on a combination of the content of the email messages and the associated email headers, does not show a statistical difference in any of the seven change-measures inspected. While this analysis was not able to demonstrate that large organizational events shift organizational communication networks, we believe that the lack of statistically significant results may be due to problems with the data which we discuss below.

Communication flow processes in the Enron corpus

Over the last six years CERI has developed a suite of methods to analyze patterns of communication flow data (who is talking to whom and when; Kiekel, Cooke, Foltz, & Shope, 2001. Kiekel, 2004). Patterns in the specific timing and sequences of interactions can be extracted using some custom routines (ProNet, Chums, Dominance). We refer to the set of routines as FAUCET (Flow Analysis of Utterance Communication Events in

Teams). FAUCET metrics have been developed and validated in a UAV command-and-control scenario and more recently in the context of Enron emails. In this project we focus on the application of ProNet (Cooke, Neville, & Rowe, 1996; Kiekel, Gorman, & Cooke, 2004).

One advantage of focusing on communication flow is that these data can be collected relatively cheaply and unobtrusively compared to content data that requires either speech recognition routines or laborious human transcription and coding. In addition, specific interaction or flow patterns have been found to be associated with effective team performance or team Situation Awareness (Cooke, Gorman, Kiekel, Foltz, & Martin, 2005). Similarly some patterns may signal a loss of team situation awareness. Although the data used in this analysis is very basic, it is also inexpensive to collect and analyze and is conducive to automation. Marrying these data with some content from the AutoMap networks could be done in cases in which one wishes to drill deeper. The combined interaction patterns and content could provide a representation of team coordination of the kind needed to share dynamic information in real time when knowledge and information are distributed across team members.

One of the metrics included in the FAUCET suite is ProNet (Cooke, Neville, & Rowe, 1996). ProNet is a method for reducing sequential event data that relies on the Pathfinder algorithm (Schvaneveldt, 1990) as its kernel. ProNet has been successfully applied to communication data (Kiekel, 2004). The result of this analysis is a graph structure in which communication events that occur together frequently are connected by directed links. ProNet, like Pathfinder, is limited in the sense that the multiple-link paths represented in the network structure are only certain to exist on a pair-by-pair basis.

ChainMaster is a software tool that implements the ProNet algorithm, but that extends it by doing tests for the existence of chains at multiple lags. With these tests, likely multiple-link paths can be highlighted. The ChainMaster analyses reported here used the following Pathfinder parameter settings: Number of nodes = 5, q = 4, and r = infinity. These parameter settings are the default values that the Pathfinder algorithm uses to generate networks from proximity matrices. ChainMaster provides regularly occurring multilink chains as output when provided with XML-formatted flow data as input. In this project, the Enron email database was examined by applying ChainMaster to segments of it to uncover regularly occurring chains. Resulting chains were then examined for changes corresponding to critical corporate events.

Data preprocessing

Using the Enron e-mail corpus described above, we categorized entities into one of nine particular job functions: president, vice president, CEO, director, manager, lawyer, trader, employee, or unknown. The database was then filtered by selecting only entities categorized as "Executive Group" members. The "Executive Group" consisted of those email entities categorized as president, vice president, CEO, director, and manager (Table 16.2). In order to reduce the number of spurious links in the ChainMaster networks due to undirected email traffic such as "list serves," the data were then filtered by selecting those emails only sent within the "Executive Group"

Table 16.2. Executive group

ChainMaster Node	**Job Category**
1	President

2	Vice President
3	CEO
4	Director
5	Manager

Segments of Enron email flow data were processed using the ChainMaster software tool using the Pathfinder parameter settings: Number of nodes = 5, q = 4, and r = infinity. ChainMaster returned regular occurring chains for each segment of emails processed.

Results

The immediate objective was to use ChainMaster to detect shifts in the email flow patterns. In order to accomplish this several analyses were conducted.

First, multiple non-critical controls were identified and compared to five critical Enron events. Of these critical events, one was excluded because stable chains were not identified. Critical event time periods were defined as the critical event (e.g., the day on which the event occurred) plus or minus three business days. Control time periods were also identified as a non-critical day plus or minus three business days. The control time period took place one week before a critical day. The original non-critical events included the dates: August 1st, September 11th, October 2nd, October 23rd, and November 20th 2001.

Critical and control chains from the ChainMaster analysis resulted in both common and uncommon chains. An uncommon chain is defined as a chain consisting of nodes and directional links (e.g., President → CEO) that occurs in either the control or critical time period, but not both. Alternatively, a common chain occurs in both control and critical time periods. In order to measure change between control and critical time periods in email flow, nodes and directed links between nodes for each time period were compared. Because ChainMaster is based on transition matrices for each time period, the difference between transition matrices for baseline versus critical were computed as the C-value, where C-value is computed as C = number common links / number unique links. A relatively high C-value indicates small change between baseline and critical chains (small change in communication patterns), while a relatively low C-value indicates large change between baseline and critical chains (large change in communication patterns).

C-values were calculated for all four critical events relative to a non-critical control. The results of this analysis indicate that the smallest C-values (i.e., the biggest difference or change in flow patterns) occurred for Events 1 and 4; respectively when Jeffery Skilling resigned and when Enron filed for bankruptcy (Figure 16.2). Therefore in this analysis the largest change in email flow was detected in the days following and leading up to these two critical events.

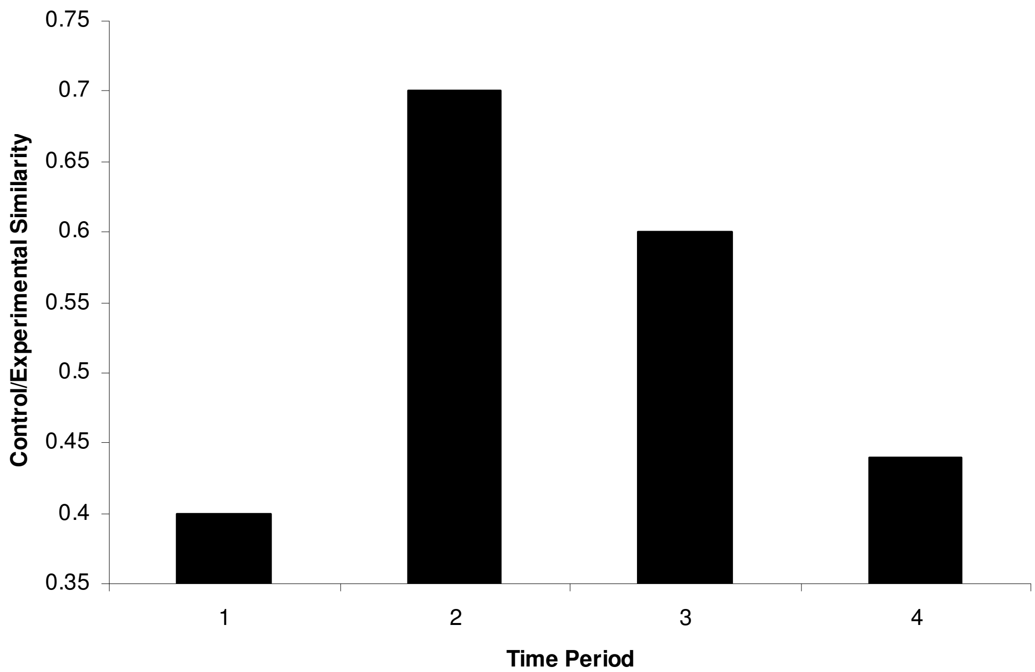

Figure 16.2 Similarity of flow patterns between critical Enron events and non-critical controls. High similarities indicate less change from baseline.

A second analysis using a single baseline was undertaken due to overlaps between non-critical control periods and critical event time periods. For this analysis, the non-overlapping August 1st control time period was chosen as the single baseline in order to detect change in the four critical time periods and a single non-critical control (Table 16.3).

Table 16.3 Dates and descriptions of critical or non-critical events

Date	**Event**	**Description**
11-Sep	Non-Critical Control	X
14-Aug	Critical Event 1	Jeffrey Skilling resigns

26-Sep	Critical Event 2	Kenneth Lay reassures Enron employees
22-Oct	Critical Event 3	Enron announces SEC investigation
2-Dec	Critical Event 4	Enron declares bankruptcy

C-values were calculated for all four critical events relative to the non-critical baseline. Further, a C-value was calculated for a non-critical control event relative to the non-critical baseline. The C-values for this analysis are presented in Figure 16.3.

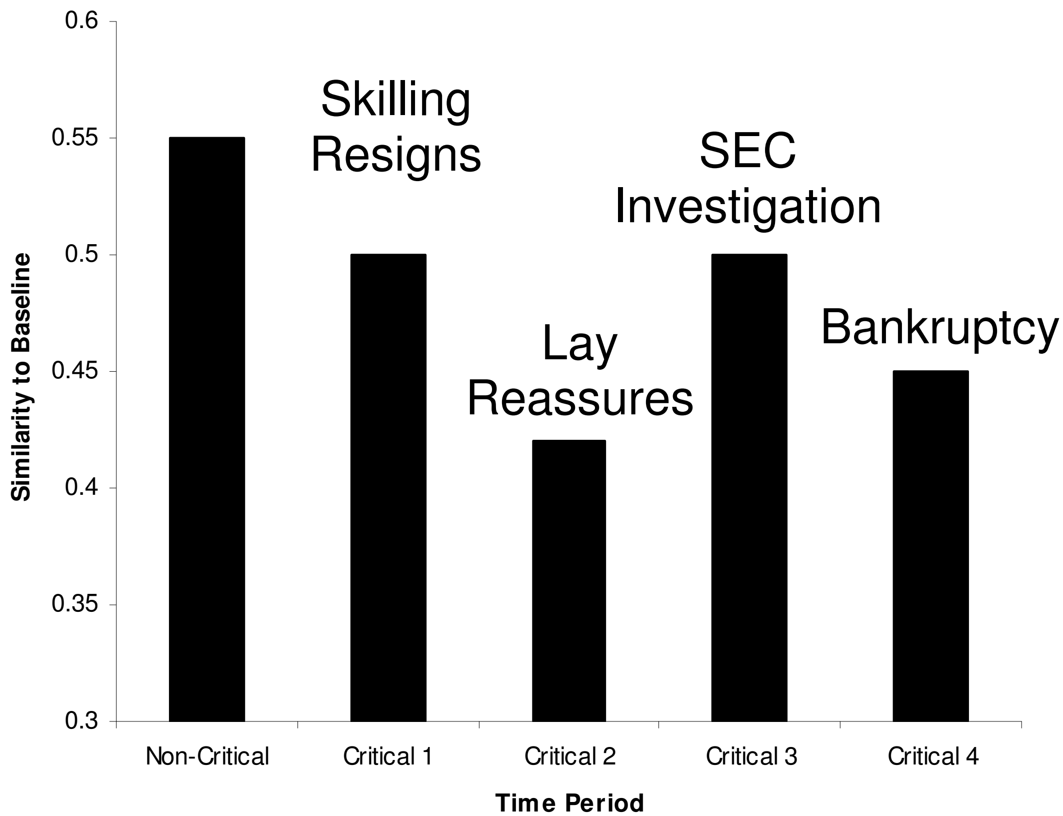

Figure 16.3 Graph of C-values for the critical versus the August 1st baseline and one non-critical event

The C-value analysis detected the degree of change in *nodes* between time periods. The biggest change occurred (i.e., lowest similarity) when Ken Lay reassured Enron employees (Critical Event 2). The relatively high C-value for the Non-Critical control event indicates the presence of consistent email flow patterns between the two non-critical event time periods.

Further analysis was undertaken in order to detect change in the chains (i.e., beyond pair-wise links) between the two time periods. The number of common chains divided by the number of unique chains between time periods was calculated in order to estimate change in chains. This analysis was again done using the single baseline time period centered on August 1st. The results are graphically depicted in Figure 16.4.

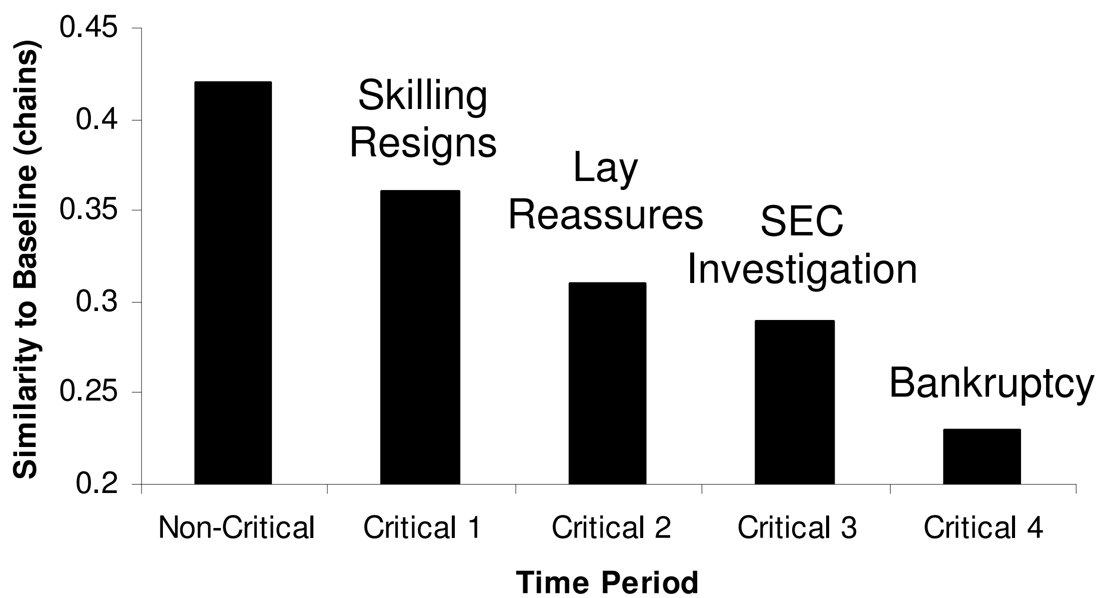

Figure 16.4 Chain difference between non-critical and baseline and chain difference between critical time periods and the August 1st baseline

The results indicate that the chains differ increasingly from the baseline for the critical events as time separation increases from baseline. However, this was not the case for the non-critical control chains. The non-critical chains were actually temporally between Critical Event period 1 and 2, suggesting the possibility that the non-critical control period and the baseline had similarities that were independent of time. Interestingly, the non-critical (in the context of Enron events) control period also included the date of September 11, 2001. This suggests that changes that are detected here in flow patterns may be specific to the Enron events. More data are needed to support this possibility.

Building on the previous analysis of change in chains, the chains that consistently differed between the critical and baseline time periods were explored. Consistent change was defined as a chain that appeared in more than half of the critical time periods but did not appear in the non-critical August 1st baseline. A chain that was consistently present during the critical events was a president to vice president pattern (4/4), and to a lesser extent, vice president to president (3/4). In the August 1st baseline chains, the president was only linked to the managers. However, in the non-critical control time period, the president to vice president and vice president to president chains were also detected. Overall this pattern suggests more communication among those at the top of the chain of command during critical events.

Discussion of flow analysis

The utility of detecting change in team or organizational dynamics, and the exact nature of this change, is in measuring SA from an interaction-based perspective. Systemic SA (Walker, Stanton, Jenkins, Salmon, Young, Beond, Sherif, Rafferty, & Ladva, 2006)) and Coordinated Awareness of Situation by Teams (CAST; Gorman, Cooke, & Winner, 2006) are two methodologies for measuring SA that have recently been developed to measure organizational and team SA, respectively. The principle goal of these approaches is not to measure the overlapping or divided situational knowledge of individuals, but rather how the team or organization as a whole changes given an evolving situation.

The ChainMaster tool, within the larger IMAGES framework, shows promise as a method for detecting organizational changes in information flow using low-level, content-free interaction data. Shifts in ChainMaster flow patterns in Enron emails seem to correspond with significant corporate events. Based on the process or interaction view, Team SA is the coordinated perception and action of team members in the face of change. Thus, coordination shifts as seen in these data would be anticipated in response to critical corporate events. These data indicate that flow shifts may be useful as a signal that an organization is adapting to environmental change. Likewise, lack of response may be indicative of a lack of SA on the part of the organization.

The application of ChainMaster to emails stretches the limits of this approach for several reasons. One limitation ChainMaster encountered in analyzing the Enron email corpus involves the detection of spurious links between aliases that were not connected in reality. This limitation is most serious in large, highly unconstrained samples, for example when list serves are used. A second limitation ChainMaster encountered was in analyzing relatively small time periods of emails (i.e., three days). Specifically, when an unreported analysis was conducted on three-day time periods ChainMaster was unable to detect email patterns in seven out of ten cases for periods that both preceded and followed the critical and non-critical dates.

The results from the ChainMaster analysis of flow patterns demonstrate a means of using communication data to measure team SA within a process-oriented framework. The results also support the validity of a process-oriented view of Team SA. That is, critical corporate events are associated with process changes in the organization. Very basic and easy to collect changes in communication flow appear promising as indices to shifts in organizational interaction in response to environmental change. The fact that change can be detected using these very low cost, basic measures, has enormous potential for on-line monitoring of organizational communications for real-time intervention.

Discussion and Conclusions

The goals of the current research effort were simple: explore ways in which the situation awareness of an organization could be derived using only the indicators of behavior resident in captured communications. If organizational SA could be assessed rapidly and largely automatically for trainers, managers, and operations, overall organizational performance might improve. These assessments would be especially valuable in complex organizations with critical responsibilities, such as the Expeditionary Strike Group or Air & Space Operations Center.

Situation awareness is not straightforward to define or measure. The three stages of situation awareness described by Endsley (1988)—which roughly correspond to perception, understanding, and prediction—are largely internal to the individual. Over the past 20 years, validated measures of situation awareness have been developed for small teams engaged in complex tactical tasks. These measures, involving during-event probes and post-event interviewing techniques, have allowed researchers to better understand the dynamics of these organizations. Unfortunately, the organizations of interest in the current research are significantly different. The timescale of their work is longer; their work is done over the course of hours, days, and weeks, not seconds and minutes is the case in aviation or surgical settings. The size of the organization is larger and thus has different knowledge requirements and interactivity patterns than small organizations of only 3 – 8 members. The sensitivity of tasks to interruption is low, and this makes some current measures of situation awareness (which require that users pause their tasks to respond to polls or probes) inappropriate for real-world use.

Clearly, new measures of situation awareness are needed for large, complex, distributed organizations. In as much as communication enables the work within and between these Joint Cognitive Systems of people and technologies, an analysis of communication can lead to insights into organizational performance, interaction, and ultimately situation awareness. The types of analyses possible on communication are many and varied, from time intensive utterance decomposition based on particular informative taxonomies to complex statistical models of word usage in a multidimensional space. Each technique gives us a glimpse into particular aspects of the inner workings of organizational behavior. When taken together, these perspectives give a more complete picture and understanding of the organization. In the present work, we have developed and integrated tools that facilitate communications analysis at the organizational level. There are other communications analysis tools and metrics that may have value in specific circumstances, and new tools are being developed in academia and industry. IMAGES was designed in a modular fashion, so that it can interoperate with other tools and metrics with relative ease.

We applied IMAGES to the Enron e-mail corpus to evaluate its potential. Although not conclusive, our analysis suggests that we may be able to identify changes in patterns of interactions or network configurations that are indicative of organizations in distress. For example, analyses of networks show critical events seem to isolate different components of an organization – as derived from communications patterns – from each other. These components increasingly focus on different topic areas. These areas may be related to critical events in their areas of expertise. Over time, this segregation abates and normal interactivity resumes.

Similarly, patterns of interactivity among communicators sometimes emerge during critical events that do not arise during more stable periods. In any environment, there is a baseline pattern of interactivity and communication among group members; a level of variability in message flow that is a natural complement to the work required in that organization. Some events disrupt this flow, causing deviations from the baseline that could eventually be identified and used in a predictive manner. Our findings on this topic were made using the FLOW analysis tools, but they are less conclusive than we would like. The medium of communication used in the Enron corpus—e-mail—is an imperfect match with FLOW analysis. Because e-mail is asynchronous and multiple conversational threads interwoven, the ordinality of the messages (the flow) became less clear and thus less informative than other media, such as voice communications and text chat.

The Enron e-mail corpus

The results of this research are promising. However, the reader should consider the limitations concerning our findings. These are due in part to the nature of the Enron corpus itself. The corpus was chosen because of its apparent similarity to our target military organizations. Its availability, size, and complexity made it widely used in the machine learning and text processing areas. However, it is critically lacking in some fundamental respects.

- **The corpus includes only part of the email from part of the organization.** The corpus was constructed by aggregating the e-mail messages found on a small number of machines, not from a central server. It was intended to assist in the court proceedings concerning a few members of Enron, and thus was not intended to comprehensively cover *all* communications among *all* employees of Enron. The intended environment for IMAGES, in contrast, is a domain in which all communications are captured and stored.
- **The corpus systematically omits communication over some channels.** The Enron corpus was collected from e-mails sent in the early 2000s. At that time, e-mail was widely used and accepted as a primary mode of communication. However, much communication at Enron occurred in face-to-face settings, over the telephone, and in written form (e.g., point-to-point facsimile, broadcast memoranda). An analysis of an e-mail corpus, even if it were complete, would lead to an imperfect picture of the organization. It is plausible that high-criticality communication occurs in more immediate communication media like face-to-face and telephone communications, leading to spurious results when the messages are analyzed.
- **The critical events lack equivalency.** To demonstrate the ability of the communications assessment techniques to detect changes in organizations, we chose five high profile events. There is a significant amount of information about these events and their effects on the company, employees, and larger business community. However, we did not make any attempt to ensure that these events were of similar type or scale. Thus, it is likely that different critical events engaged different employees and elicited different reactions from them, leading to different patterns of interactivity and different topic profiles. Future analyses will be enhanced by categorizing critical events and theorizing on the likely effects of those events on communication.

Given these caveats, it is encouraging that we found persuasive trends in these data. Specifically, we found that communications differed between date ranges in which there were critical events, and periods in which there were none (control conditions). This demonstrates the robustness of the approaches used. Further research is needed to explore the full potential of these approaches in conditions that more exactly map to the environments of interest.

The techniques we demonstrated on the Enron corpus can be employed on other large and complex organizations. Given a baseline understanding of the communications characteristics of an organization, users of IMAGES can detect changes in those characteristics and assess deviation from normalcy. When Air Force personnel in an Air and Space Operations Center are being trained using a known scenario, we can compare their communications patterns with previous successful and unsuccessful teams. Both content and flow measures could be used as diagnostic tools, feeding back information to trainers and evaluators. In less predictable environments, such as a deployed Expeditionary Strike Group, the challenge is to discriminate between the normal and abnormal changes in interactivity and communication patterns. This will require additional research. In the interim, gains in operational efficiency may be made simply by presenting changes to operations officers who have an intuitive feeling for congruous and incongruous patterns given the mission changes.

Situation Awareness and Macrocognition

We made two fundamental arguments at the beginning of this research program. First, we argued that organizational level situation awareness differs fundamentally from what is commonly accepted as situation awareness in tactical environments. As a knowledge construct, organizational SA is the allocation of knowledge and understanding around an organization, with some information shared and some unshared between roles (which have different task demands) and over time. Organizational situation awareness is, thus, a complex mix of shared and team situation awareness. As a process construct, organizational SA consists of patterns (flows) of communications that enable that knowledge to be shared and acted upon in a constantly changing, dynamic environment. It is only when both the knowledge and process aspects of team cognition are considered that an indication of organization situation awareness can be derived.

The second argument fundamental to this research was that organizational situation awareness can be derived from analysis of communications patterns. Teams do not work silently; communication enables information to be passed, confirmations to be made, and ambiguities clarified. An analysis of the externalized communications of a Joint Cognitive System is a window into cognitive collaboration. We have argued that communication *is* team cognition, or at least that it both enables and reflects team cognition to a high degree. We have taken a broad perspective on approaches to communications, including measures of content, context, and process. Given the definition of organizational situation awareness, all three of these aspects of communications need to be considered.

For this operationalization of situation awareness, the communications measures we propose are very promising. Even with a sub-optimal corpus, trends implied that changes in communications patterns were related to critical changes in the environment. These trends appeared both in context and process measures of communications. While considerable research remains to refine these measures, the potential is clear: analysis of communications offers a window into organizational dynamics and situation awareness.

These same measures of communications have implications for a theory of macrocognition. Communications is an enabler of macrocognitive processes of knowledge interoperability development and shared understanding (Warner et al. 2005). Clearly, teams of individuals in problem solving or decision making situations must communicate to succeed in their tasks, to gain common ground, and to manage ambiguities. Analysis of communications can thus provide insight into critical team processes.

The macrocognition theory of Warner, et al. (2005) was developed primarily to explain the behaviors of teams. The analyses performed for this project concerned a very large organization, one several orders of magnitude greater in size and perhaps in task complexity that teams. In organizations (unlike teams), consensus is not a goal and multiple macrocognitive processes may occur simultaneously in different parts of the organization. Thus, future research should test the processes and stages of the Warner et al (2005) theory of macrocognition in larger organizations that contain and support the operational teams that accomplish critical missions.

Into the future

The work described in this report represents the first steps on a long journey. It is only in the past decade that communications have been largely capturable and storable, and only in the past few years that researchers have considered analysis of communications a viable approach to performance assessment. In this study, we applied multiple, converging approaches that had never been combined, and used them on a corpus that was orders of magnitude larger and more complex than had been previously attempted for this purpose. The results are promising, with trends in the directions of interest.

As we proceed, those trends need to be investigated further. In an environment in which there is greater control and understanding of the organization and its tasks, we could more accurately correlate changes in communications behaviors with external factors. A large-scale training session or exercise would be an ideal venue in which to apply IMAGES, as the scenario would be carefully scripted and the participants identified. In such an environment, existing measures of performance could be correlated with patterns in communication, allowing a stronger understanding of the antecedents of those patterns.

As discussed above, the Enron corpus – while a useful data set in some respects – was a small sample of the communications and communicators at Enron. This paucity of communications led us to interpret patterns in the data conservatively, as trends. In the future, the same communications measures should be applied to a more complete corpora of communications. An ideal corpus would capture all of the communications, not just those that occur in e-mail form. Unobtrusive methods to capture text chat, recorded/transcribed telephone, and face-to-face communication are all being developed. This richer, captured data would be a good proving ground for the communications measures used here. The choice of communications media itself may also be an indicator of team performance, and media measures could be used to complement or confirm measures used here.

Organizations differ in purpose, composition, complexity, and a host of other dimensions. As this research proceeds, the individual characteristics of organizations have to be considered as they relate to the analysis of their communications. How do the communications patterns of an Expeditionary Strike Group differ from those of a similarly sized Army division? How does it differ from a corporate entity? Understanding these differences will help us refine our interpretation of patterns and their implications.

People have been communicating for millennia; we have only begun to understand how and why. The measures, methods, and tools developed in the IMAGES project should improve both our knowledge of communications and our management of teams, whose missions hinge on accurate, effective communication.

Acknowledgments

Work described in this chapter was supported by N00014-05-C-0505 from the Office of Naval Research. The views, opinions, and findings contained in this chapter are the author’s and should not be construed as reflecting the official views of the Department of Defense, Aptima, Inc., or Carnegie Mellon University.

References

Alberts, D.S., (2002), *Information Age Transformation: Getting to a 21st Century Military*. Washington, DC, CCRP Publications. First published 1996.

Carley, Kathleen & Diesner, Jana & De Reno, Matt. (2006). AutoMap User's Guide. *Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISRI-06-114.*

Cooke, N.J., Gorman, J.C., & Keikel, P. (2007). Communication as Team-level Cognitive Processing. *Macrocognition in Teams.* Letsky, M., Warner, N., Fiore, S., & Smith, C., Editors. Ashgate: United Kingdom.

Cooke, N.J., & Neville, K. J., & Rowe, A. L. (1996) Procedural network representations of sequential data. *Human-Computer Interaction, 11,* 29-68.

Diesner, J; Frantz, T & Carley, K. (2005) Communication Networks from the Enron Email Corpus “It's Always About the People. Enron is no Different”. Computational & Mathematical Organization Theory, 11(3):201-228

Endsley, M. R.(1988). Situation awareness global assessment technique (SAGAT). Proceedings of the National Aerospace and Electronics Conference (NAECON). (New York: IEEE), 789-795.

Gorman, J. C., Cooke, N. J., Pedersen, H. P., Connor, O. O., and Dejoode, J. A. (2005). Coordinated awareness of situation by teams (CAST): Measuring team situation awareness of a communication glitch. *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting* (pp. 274-277). Santa Monica, CA: HFES.

Gorman, J.C.., Cooke, N. J., & Winner, J.L. (2006). Measuring team situation awareness in decentralized command and control systems. *Ergonomics, 49,* 1312-1325.

Kiekel, P.A. (2004). Developing automatic measures of team cognition using communication data. Ph.D. Thesis, New Mexico State University.

Kiekel, P. A., Cooke, N.J., Foltz, P.W., Gorman, J. C., & Martin, M.J. (2002). Some promising results of communication-based automatic measures of team cognition. *Proceedings of the Human Factors and Ergonomics Society 46th Annual Meeting*, 298-302.

Kiekel, P. A., Cooke, N. J., Foltz, P. W., & Shope, S. M. (2001). Automating measurement of team cognition through analysis of communication data. In M. J. Smith, G. Salvendy, D. Harris, and R. J. Koubek (Eds.*), Usability Evaluation and Interface Design*, pp. 1382-1386, Mahwah, NJ: Lawrence Erlbaum Associates.

Kiekel, P.A., Gorman, J. C., & Cooke, N. J. (2004). Measuring speech flow of co-located and distributed command and control teams during a communication channel glitch. *Proceedings of the Human Factors and Ergonomics Society 48th Annual Meeting*.

Quek, F., McNeill, D., Bryll, R. Duncan, S. Ma, X., Kirbas, C., McCullough, K. E., Ansari, R. (2002). Multimodal human discourse: Gesture and Speech. ACM Transactions on Computer-Human Interaction (TOCHI), 9, 171-193.

Schvaneveldt, R. W. (1990). *Pathfinder associative networks: Studies in knowledge organization.* Norwood, NJ: Ablex.

Warner, N., Letsky, M., Wroblewski., E. (2007). Empirical Model of Team Collaboration: Focus on Macro-Cognition. *Macrocognition in Teams.* Letsky, M., Warner, N., Fiore, S., & Smith, C., Editors. Ashgate: United Kingdom.

Weil, S. A., Carley, K. M., Diesner, J., Freeman, J., & Cooke, N. J. (2006). Measuring Situational Awareness through Analysis of Communications: A Preliminary Exercise. Command and Control Research and Technology Symposium, San Diego, CA.