

A Multi-Agent View of Strategic Planning Using Group Support Systems and Artificial Intelligence

R. ORWIG

University of Arizona MIS Department

H. CHEN

University of Arizona MIS Department

D. VOGEL

University of Arizona MIS Department

J. F. NUNAMAKER, JR.

University of Arizona MIS Department

Abstract

The strategic planning process is dynamic and complex. Including a Group Support System (GSS) in the problem-solving process can improve the content quality of the strategic plan by allowing increased participation by more members of the organization. However, it can also add to the complexity of the problem by increasing the quantity of textual information that can result from group activity. Added complexity increases cognitive overload and frustrations of those participants negotiating the contents of the strategic plan. This article takes a multi-agent view of the strategic planning process. It considers group participants as multiple agents concerned with the content quality of the strategic plan. The facilitator agent is responsible for guiding groups in the strategic plan construction process as well as for solving process problems such as cognitive overload. We introduce an AI Concept Categorizer agent, a software tool that supports the facilitator in addressing the *process* problem of cognitive overload associated with convergent group activities by synthesizing group textual output into conceptual clusters. The implementation of this tool reduces frustrations which groups encounter in the process of classifying textual output and provides more time for discussion of the concepts themselves. Because of the large amount of convergent activity necessary for strategic planning, the addition of the AI Concept Categorizer to the strategic planning process should increase the quality of the strategic plan and the buy-in of the participants in the strategic planning process.

Key words: strategic planning, group systems support, electronic meeting systems, concept clustering, group problem solving, facilitation of group processes

1. Introduction

As organizations seek to adapt in a world of rapid change, strategic planning becomes increasingly dynamic and complex. Group Support Systems (GSS) provide a means by which a larger number of organizational stakeholders can efficiently and effectively participate in the planning process. The resulting plan should benefit by the richness of knowledge provided by the greater representation of organizational members and should

be more easily implementable due to the greater “buy-in” that should result from increased participation. A success factor critical to this involvement is the successful organization of massive amounts of information generated by such a group.

A variety of ideas go into any given strategic plan: mission statements, vision statements, goals, objectives, critical success factors, business processes, value chains, etc. Experience with GSS has shown a recurring pattern of use in (1) generating, (2) organizing, and, perhaps, (3) prioritizing these ideas. We refer to this recurrent pattern as divergence, convergence, and consensus polling, respectively. In most instances, it is much easier to generate ideas than to organize them effectively and efficiently. If idea organization is delegated to an individual or small group, there exists a high risk of not identifying important clusters of ideas and losing group buy-in to these structured ideas. If the organizing task is done by the group as a whole, the amount of time consumed can become excessive, with group members becoming frustrated and “dropping out” of the process.

The idea organization (convergent) process contributes to the creation of what Lewis and Sycara (1993) call the “shared model”: “a more comprehensive model of the problem at hand...which incorporates elements of others’ [agents’] expertise. The ‘shared model’ between two agents also defines a common vocabulary that the two agents can use to communicate in an intelligible way.” Each participant in the strategic planning process can be viewed as an agent contributing viewpoints and expertise toward creation of one strategic plan.

The responsibility of managing any strategic planning process (including the subprocesses of idea organization) is typically put upon a supervisory agent. We call this agent the facilitator. We view the participants as multiple agents responsible for creating the *content* of the strategic plan. The facilitator is the agent responsible for managing the plan creation *process*. Because of the problems, frustrations, and great amount of time experienced in idea organization, we introduce an artificially intelligent agent to provide concept categorization support to the facilitator and group responsible for developing the strategic plan.

Artificial intelligence (AI) has been shown to be successful in clustering ideas generated by a group (Chen 1994). Questions arise again, though, regarding the degree of stakeholder “buy-in” and overall effectiveness as well as efficiency of the AI approach in operational contexts. AI, in combination with facilitation and group participation in the idea organization exercise, however, shows excellent promise.

In this article we introduce and describe strategic planning in terms of the components of the process of creating the plan. We then use an example to describe strategic planning in the GSS setting and explore the application of AI-assisted categorization in concert with facilitation and group participation. The use of AI-assisted categorization is compared to unassisted facilitator categorization as well as to interactive and interactive use in organizational sessions. Initial results from the use of the prototype AI Categorization tool are discussed. We conclude by addressing issues for future research with respect to the dynamic use of AI-assisted categorization in GSS strategic planning context.

2. Strategic planning

Strategic plans contain a variety of elements. Strategic planning typically involves establishing or validating an organizational vision and/or mission statement, identifying corporate values, goals and objectives, and perhaps specifying critical success factors or tactics in support of the goals and objectives. It may also include identification of processes and value chain links or business processes for business process reengineering. The strategic planning process may include organizational identification exercises such as Strengths, Weaknesses, Opportunities, and Threats (SWOT) Analysis. Definitions for these terms vary across organizations, and organizations consider a variety of components as essential elements of their strategic plans.

Figure 1 provides a representation of a goal-directed approach to the construction of a strategic plan. The Strategic Plan is the final goal represented by the box to the extreme right of the figure. The components of the strategic plan are represented as subgoals. Subgoals may or may not be further decomposed, say, into sub-subgoals. The actual decomposition of a specific organizational strategic plan may not be symmetric, as is depicted in Figure 1. For example, a SWOT Analysis may be further decomposed into meetings that identify strengths, weaknesses, opportunities, and threats separately. The strategic planning process involves managing the construction of plan components. This includes determining which components need to be part of the strategic plan, matching the

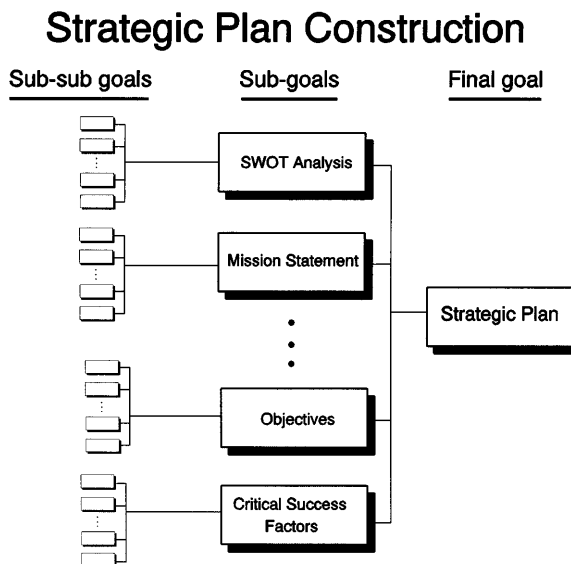


Figure 1. Goal-oriented view of strategic planning.

appropriate participants from the organization with each specific goal, and monitoring progress. It also includes the information management tasks of collection, synthesis, and validation of all ideas, topics, and comments that go into the components of the plan.

Another approach to viewing the strategic planning process is that of managing inter-related problems. Ackoff (1981) uses the term “interactive planning” to conceptualize a participative approach to dealing with a set of interrelated problems. Through interactive planning, recognition of the past, sensitivity to the present, and foresight towards the future are seen as different but inseparable aspects of the planning process. The objective is to focus equally on the past, present, and future in a participative fashion involving relevant stakeholders at each point along the way.

Ackoff describes three operating principles associated with interactive planning: the participative principle, the principle of continuity, and the holistic principle.

- The participative principle asserts that the principle benefit from planning comes from engaging in it, and that, in planning, process is the most important product. Executives and organizational members at all levels need to engage directly in the planning process.
- The principle of continuity stresses that planning is a continuous process of monitoring, evaluation, and modification, as circumstances and values change, both internally and externally to the organization.
- The holistic principle has two parts: the principle of coordination and the principle of integration. The principle of coordination states that “no part of an organization can be planned for effectively if it is planned for independently of any other unit at the same level.” The principle of integration states that “planning done independently at any level of a system cannot be as effective as planning carried out interdependently at all levels.”

Maximum participation by organizational members benefits organizational strategic planning benefits in at least two ways: first, increased participation increases the richness of information that can supply novel approaches to complex planning problems. Second, large group participation in the strategic planning process should increase buy-in of the resulting plan. Participants have a greater stake in the strategic plan and can better communicate the strengths of the plan to those organizational members who did not physically participate in the creation of the plan. Thus, a greater number of organizational members participating in the strategic planning process logically leads to a more complete and implementable plan. Participants serve as expert agents for their knowledge areas and/or for their respective functional areas. However, the very strengths of varieties of expertise and viewpoints among the agents which contribute to a quality strategic plan become problems during the strategic plan construction process. Working the large number of ideas and viewpoints into a single strategic plan is a complex process. This complex process involves the negotiation of the various views and expertise into one shared model (Lewis 1993). As with most complex negotiation situations, an outside agent is often used to support the negotiation process.

We define a facilitator agent as that agent responsible for the overall strategic planning process. This includes managing the complex negotiation processes that are required

among those participants collaborating on plan components. The facilitator is also responsible for designing and managing the group processes that result in creation of the products that satisfy the goals of the plan components represented in Figure 1.

Experience with facilitating strategic planning groups has shown that a recurring pattern of three stages occurs in the construction of each plan component. These three process stages are: group divergent activity, convergent activity, and, optionally, a consensus polling or process check.

Divergence. For example, assume a subgoal of a strategic plan is a mission statement. Also assume that 20 or more people are gathered together as a group to create the statement. Depending upon the characteristics of the group, at least three methods of creation are possible: (1) have each person write a mission statement—thus creating 20+ “straw-men” statements; (2) modify one “straw-man” mission statement prepared in advance of the meeting; and (3) brainstorm on what elements need to be in the mission statement, obtain group consensus as to what those elements need to be, and use those elements to build the mission statement.

Divergence in a group setting means that members of the group provide comments or ideas concerning the topic at hand. If 20 or more people each write a separate mission statement, a large quantity of mostly redundant text will be generated. If a “straw-man” mission statement is offered, it is usually appropriate to allow group members to discuss the sample mission statement in order to add missing elements or to suggest changes in the wording. Here, again, group members might supply a generous quantity of comments. The “straw-man” mission statement may cause a “groupthink” mission statement that inhibits the introduction of new elements into the mission statement. If fresh thinking and a totally new mission statement is desired the group might build the mission statement from scratch in a bottom-up manner. Group members may suggest ideas for elements such as “customer satisfaction,” “profitability,” or “long-term growth” in an electronic brainstorming session (a divergent activity). Again, twenty or more people might suggest a large number of elements in a short period of time. The group must identify these elements and converge on a list of essential elements for the Mission Statement. After voting on which elements are essential, the mission statement could then be written using the most important elements.

Whatever the method of building the mission statement, the generation of ideas and suggestions from twenty or more people typically results in hundreds of lines of text.

Divergence allows every member of the group opportunities to contribute to the subgoal and, ultimately, to the output of the meeting. Facilitators will choose the method of divergence that they feel will be most appropriate to the characteristics of the group, task, and context of the meeting. Whatever the method, participation by group members in the divergent stage typically creates a large volume of text that needs to be synthesized in some way.

Convergence. Convergence refers to the synthesis of the text that is generated from the divergent stage of the problem-solving process. A common objective of the convergence stage is to produce a list or outline of topics that address the current subgoal of the

electronic meeting. The intent is to categorize the comments or ideas found in the text into a topical list or list of concepts. In the corporate mission example above, the objective might be to derive a list of elements that must be evident in a mission statement. In order to develop this list of elements, group members need to browse the comments produced during the divergent stage of any of the three methods described.

Several facilitation techniques can be used to address some of these convergence problems. For instance, with respect to achieving an appropriate level of abstraction, a variation of the Nominal Group Technique might be used whereby participants can be asked to propose one topic each. One suggested topic can be collected from each group member and posted on a public screen. The group can then browse this list on the public screen and evaluate the items with respect to relevance to the subgoal and level of abstraction. The facilitator guides group members in looking for duplicate entries and topics that are too general or too specific with respect to the current goal. If the group continues to have trouble settling in on a level of abstraction, another round of one-item suggestions can be collected and evaluated.

Duplicate topic items or synonymous topics can be merged together under the topic heading that seems most appropriate to the group. Items that are too general or inappropriate are typically deleted. For example, in the issues contained within a mission statement, a participant may suggest “maintain business in a legal manner” as one of the elements that need to be contained. This may be a topic that is simply understood and too general to be used in a given organization’s mission statement.

Polling/consensus check. The polling/consensus checking stage of the problem-solving process involves sampling members of the group to determine how much agreement exists concerning the topics that were created on the list created during the convergent stage. It can also serve as an evaluation of how unified members of the group are with respect to their understanding of the concepts.

Even if some method of polling or consensus checking is not required for the items generated from the convergent stage, meeting facilitators must check that the output list conforms to the requirements of the subgoal.

If a given strategic plan contains many components, the sequence of divergence, convergence, and polling stages will occur many times during the group strategic planning process. Experience has shown that the convergent stage is the most problematic. Groups usually have no trouble expressing ideas or commenting during divergence. However, as mentioned above, the convergent stage is really a negotiation stage. It is during convergence that diverse opinions, viewpoints, and functionalities must come together to create a shared plan component, such as a list of corporate goals or objectives. Embedded within the items on the shared list are shared definitions and understandings of what each item on the list means. Facilitating convergence with a group of people is a negotiation problem that involves management of information and group behavior. The information management problem results from the large quantities of information generated from the divergent stage of the process. The group behavior management problem results from the variety of personalities involved in the process and the need to create one common product.

3. Strategic planning in GSS

Group Support Systems (GSS) provide a means by which a large number of organizational stakeholders can efficiently and effectively participate in the planning process (Nunamaker 1991b). A GSS helps manage many of the problems and complexities that arise in meetings with larger groups. For example, participants in the meeting can all “speak” at the same time in discussing issues through individual PC workstations. This helps manage the slow, serial nature of non-GSS meetings. Comments and ideas can be submitted anonymously. This reduces complexities associated with matching ideas with the person providing the idea. Ideas and comments can be considered separately from the personalities who provide them. All comments are stored within the system, and can always be revisited. Participants may reread comments in light of newly made comments, or to clarify an understanding. This aids in the information management problems associated with manual marking boards, paper, and notes.

The strengths of the GSS—parallel activity, anonymity, and electronic recording—affect the general processes associated with construction of strategic plan components in different ways:

Divergence. GSS provides the support for a larger number of meeting participants. The ability to enter ideas into the system in parallel produces a large amount of information in a short period of time. Several electronic “conversations” can occur simultaneously. Greater numbers of meeting participants provide for a greater variety of ideas, comments, and topics. Since comments can be anonymous, participants should not feel constrained in providing ideas that they may not provide if they were to do so verbally. This also adds to the quantity of information that can be generated during the divergent stage. Since all ideas are recorded electronically, participants are free to browse all comments and make new comments wherever and whenever they wish. Thus, GSS provides the ability to create a larger and richer pool of information from which to make plans and/or decisions. It also contributes to the information management problem of the facilitator.

Convergence. It is this convergent stage that causes the greatest frustration with respect to group problem solving (Chen 1994). Parallel entry by participants in electronic meetings coupled with anonymity characteristics provide an environment for generating large quantities of text. GSS tools provide support for browsing the text in a serial manner. However, synthesizing the complete set of textual comments into a list of categorical topics or outline of concepts that are contained in the text is performed in a bottom-up manner. Coordinating this effort among large groups is problematic. The GSS tools provide software that allows each group member browsing capabilities plus the ability to post topics to a public screen. However, groups have the same difficulties associated with non-GSS meetings, as well as some new ones. Difficulties include: (1) determining the appropriate levels of abstraction of terms to be used in the synthesized topics; (2) maintaining consistency in the levels of abstraction; (3) not duplicating entries already on a list; (4) supporting topic suggestions with links to the ideas or comments that suggested that topic; (5) making sure all of the underlying comments have been reviewed for potential synthesis

and addition to the list; and (6) taking the time necessary to go through all of the comments. Our experience in facilitating groups has shown that groups typically cannot stay focused on a particular task for much longer than 40 minutes. For a divergent stage of 40 minutes the convergent stage may take 40 minutes to an hour.

Since hundreds of comments can be generated by a group in one divergent stage of a GSS session, the information management problem normally associated with the convergent stage is made much more serious. Current technology presents the hundreds of comments in one big file. Like a blind person learning about an elephant, participants may browse sections of the large set of comments without being able to step back and view the whole thing. Various facilitation techniques are used to address this problem: facilitators sometimes begin identifying key words being used during the divergent stage, and may use key-word-in-context software to begin indexing comments. Idea organization software may distribute the complete set of comments to all participants, but start each participant at a different location in the file. In this way, each person can browse and synthesize a different portion of the file to get a complete analysis of all of the comments.

GSS also provides the ability to link the comments created during the divergent stage plus any new comments to a topic heading or list item. For example, if a corporate goal of “increase market share” was synthesized from comments made by participants during the divergent stage, comments pertaining to market share can be linked to this topic heading for support and/or definition of the corporate goal. Participants are also free to browse the topic lists and make further anonymous comments pertaining to the topics.

In short, the convergent stage of the GSS session consists of browsing the output of the divergent stage to synthesize the ideas into a list of topics or concepts that address the current goal of the overall planning process. It also consists of electronic and verbal discussion/negotiation concerning what topics and concepts are appropriate for the list. Inherent in this discussion/negotiation is the creation of the common view of what the topics and concepts mean to the group as a whole. This discussion/negotiation is an important component to the buy-in of the group to the resulting list of the convergent stage.

Experience with GSS finds that group satisfaction levels decrease sharply when confronted with the information management problem associated with the convergent stage of the process. Facilitators must quickly determine how best to get the group to some common ground of agreement as to what items should be on the list. Starting with a list of keywords and linked comments is a start, but the key word list may be semantically skewed due to the facilitator’s ability to browse the comments as they are created or due to the facilitator’s limited content knowledge.

Consensus polling. GSS provides much additional support for determining consensus or buy-in through voting tools that allow for ranking, weighing, or otherwise evaluating list items. Output of the voting tools allow for analysis of the spread of the votes in the group—again, anonymously representing the votes. The standard deviation and distributions of the collected votes serve as indicators of levels of group agreement. Displays of the GSS can provide identification of those concepts needing further discussion and

clarification. The displays can also help identify specific issues or concepts that members of the group understand but for which negotiation may be needed in order to provide consensus in evaluation.

The convergent stage of any group task consists of overcoming the information management problems associated with the large quantities of information created during the divergent stage, as well as the negotiation problems associated with facilitating the creation of a common set of concepts, definitions, or model among the large number of meeting participants. The information management problem is a process problem in that it is unrelated to the content of the information contained in the meeting. The information management problem merely serves as an impediment to getting to the point to be able to work on the negotiation problem. We propose the introduction of a new agent, the AI Categorizer Tool, to aid in addressing this information management problem.

4. AI approach: A new agent

Participants in a GSS meeting can be viewed as a set of agents involved in creating a common model or achieve a subgoal to a group problem (creating a strategic plan). These participant agents are involved with the content knowledge of the particular group problem at hand. We have viewed the facilitator as an outside agent responsible for the *process* that the participant agents use to solve a given problem. The problems associated with the strategic planning process include many instances of potential information overload and negotiation management. Information overload problems decrease satisfaction among participants and inhibit negotiations that are essential to foster group buy-in of the meeting product. Figure 2 presents our model that introduces a third major component, the AI Concept Categorizer, as an agent to aid in overcoming the information management problems related to the convergent stage of any group problem process.

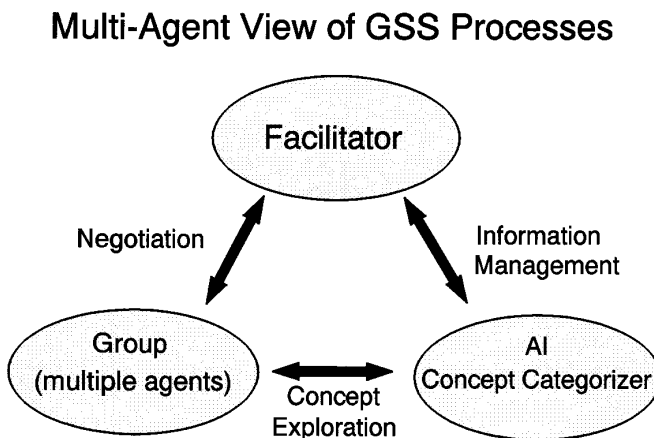


Figure 2. A Multi-agent model for convergent support.

The AI Concept Categorizer is a tool that can be used by both the facilitator and participants to address the information management problem that results from group divergent activities. It is intended to address the information management problem of GSS strategic planning for the facilitator and group by providing a straw-man list of topics or concepts discussed during the divergent stage. This straw-man list is generated by doing a textual and statistical analysis of the complete set of comments. The list is objective in that it is based on the analysis of the textual contents of the large file of divergent output. It is complete in that the full file is analyzed by one agent. This removes some of the bias that can occur when humans cannot view the divergent file as a whole. The AI Concept Categorizer divides the large file into conceptual chunks that can then be analyzed by participants and used as a starting point of negotiation. Experience with the tool in actual practice has been very encouraging and has raised many interesting issues.

Based on the characteristics of GSS meetings and our experience with several general-purpose AI and information-retrieval-based algorithms, we developed the AI Concept Categorizer tool. We adopted techniques from the following three areas: *automatic indexing*, *cluster analysis*, and *Hopfield net classification* in Chen (1994). We review these techniques in the context of our research below.

4.1. Automatic indexing

The first task for AI concept categorization is to identify the content of each individual comment. A domain-independent and computationally efficient method for content identification is the automatic indexing method, often used in information science for indexing literature. Salton (1989) presents a blueprint for performing automatic indexing, which typically includes dictionary look-up, stop-wording, word stemming, and term-phrase formation. During automatic indexing, the number of occurrences of each term in an entire divergent session is computed. By setting a document frequency threshold, we could remove some incidental and noisy terms produced as a result of the automatic indexing process.

4.2. Cluster analysis

While automatic indexing identifies subject descriptors in a comment, the relative importance of each descriptor for representing the content of the document may vary. Salton's Vector Space Model associates with each descriptor a weight to represent its descriptive power. Among the many probabilistic techniques that have been developed by various information science researchers, techniques which typically incorporate term frequency and inverse document frequency have been found to be quite useful (Salton 1989). The basic rationale underlying these two measures are: terms which appear more times in a document should be assigned higher weights (term frequency), and terms which appear in fewer documents in the whole database (the more specific terms) should have higher weights (inverse document frequency).

Based on cluster analysis (Everitt 1980), the Vector Space Model has been extended for automatic thesaurus generation (or automatic knowledge-base generation). The first stage in many cluster analyses is to convert the raw data (e.g., indexes and weights) into a matrix of inter-individual similarity, dissimilarity, or distance measures. The result of a cluster analysis will be a number of groups, clusters, types, or classes of individuals (Everitt 1980; Jones 1991; Michalski 1983a). In automatic thesaurus generation (Jones 1991; Crouch 1990; Chen 1992), the most commonly used algorithms compute probabilities of indexes co-occurring in all documents of a database (sometimes referred to as the co-occurrence analysis).

Based on an asymmetric Cluster Function we developed earlier (Chen 1992), we generated a network-like concept space of terms and their weighted relationships as follows:

$$ClusterWeight(T_j, T_k) = \frac{\sum_{i=1}^n d_{ijk}}{\sum_{i=1}^n d_{ij}} \quad (1)$$

$$ClusterWeight(T_k, T_j) = \frac{\sum_{i=1}^n d_{ijk}}{\sum_{i=1}^n d_{ik}} \quad (2)$$

Equations (1) and (2) compute the similarity weights from term T_j to term T_k (the first equation) and from term T_k to term T_j (the second equation), where d_{ij} represents the combined weight of descriptor T_j in document i , and d_{ijk} represents the combined weight of both descriptors T_j and T_k in document i . d_{ijk} is computed by:

$$d_{ijk} = tf_{ijk} \times \log df_{jk} \quad (3)$$

where tf_{ijk} represents the number of occurrences of both term j and term k in document i and df_{jk} represents the number of documents in a collection of n documents in which both term j and term k occur.

4.3. Hopfield net classification

When pairwise similarities are obtained between all term pairs, a hierarchical agglomerative cluster generation process is often adopted (Salton 1989). The cluster generation process aims to group (link) similar term pairs together. Several alternatives for linking similar terms exist: single-link clustering and complete-link clustering (Jones 1971; Salton 1978). While these information-retrieval based clustering methods have demonstrated their usefulness in clustering documents, a somewhat newer approach based on the con-

nectionist paradigm, or neural network computing, has attracted a resurgence of interest (Knight 1990). After a review of numerous networks (Lippmann 1987; Simpson 1990; Knight 1990) and based on our experience and experimentation with several major networks (Chen 1993), the Hopfield net (Hopfield 1982), which was introduced as a neural network that can be used as a content-addressable memory, appeared to be a natural candidate for divergent ideas classification. Knowledge and information can be stored in single-layered, interconnected neurons (nodes) and weighted synapses (links), and can be retrieved based on the Hopfield network's parallel relaxation and convergence methods. The Hopfield net has been used successfully in the past few years in such applications as image classification, character recognition, and robotics (Tank 1987; Knight 1990), and we have also adopted it for concept-based information retrieval in Chen (1993).

Each term in the network was treated as a neuron and the asymmetric weight between any two terms was taken as the unidirectional, weighted connection between neurons. Using each term as an individual input pattern, the Hopfield algorithm activated its neighbors, combined weights from all associated neighbors, and repeated this process for a number of iterations. The Hopfield net algorithm relied on an activate and iterate process, where

$$\text{for term } j, 0 \leq j \leq n - 1, \mu_j(t + 1) = f_s \left[\sum_{i=0}^{n-1} t_{ij} \mu_i(t) \right] \quad (4)$$

$\mu_i(t + 1)$ is the activation value of neuron (term) j at iteration $t + 1$, t_{ij} is the co-occurrence weight from neuron i to neuron j , and f_s is the continuous SIGMOID transformation function which normalizes any given value to a value between 0 and 1 (Knight 1990; Dalton 1991). This formula shows the parallel relaxation property of the Hopfield net. The above process is repeated until it reaches a fixed number of iterations or until it converges.

The resulting output revealed all other concepts that were semantically relevant to the input term. Repeating this process for all the terms according to their decreasing occurrence frequencies, we were able to identify the underlying, related concepts in the comments. The output of Hopfield net classification is thus a list of concepts, each concept being represented by a term or set of related terms that captures the semantics of the concept.

The benefits of incorporating the AI Concept Categorizer agent into the GSS convergent process can be understood best by an example. What follows is an example of an actual GSS session as viewed from the perspectives of the three agents.

5. Convergent process from three agent perspectives

A group of twenty participants were attending a Winter Meeting of GroupSystems users. These participants included managers and users of the GSS called GroupSystems (marketed by Ventana Corporation). These twenty participants were asked to use the electronic

brainstorming tool to respond to this question: *What are the most important information technology problems with respect to Collaborative Systems to be solved over the next five years?*

The group generated 201 comments over a 30-minute period. Content of the comments ranged from one-word responses to the question up to five lines of text. A sample of the comments follows (typos and misspellings are normal in the output):

Electronic Brainstorming Report

Session: technology Date: 11/18/1993 2:08pm

What are the most important information technology problems with respect to Collaborative Systems to be solved over the next five years?

1.1 development of "standard" user interfaces for remote access to info and communication

1.2 Doesn't XYZ Software give us such a standard?

1.3 XYZ is one implementation, but it certainly is not yet a standard. !!!! Maybe we need to have a standard that is NOT owned by a company which has a history of initiating look and feel lawsuits when someone else copies their "standard" !!!!!

1.4 XYZ is not a good standard even if they are the industry leader at this time. XYZ is a terrible architecture.

2.1 Voice recognition is one of the promising technologies.

2.2 The technology has been developed steadily, however, the sociological issues need more attention now.

2.3 companies that develop collaborative systems need to realize that they are working in a collaborative system!

2.4 Reliable and flexible repository systems linked to different presentation tools.

...

29.6 cellular technology hold great promise for having remote collaborative meetings but the reliability and band with need to be improved.

29.7 Bigger windows on the creations of groups... I can't see enough of what the group is doing here...

29.8 network support is going to be essential

30.1 Network bandwidth is a real problem as we get into remote screen sharing and graphics. Also, processor speeds are going to have to increase vastly to support the gesturing and other body-language-replacement features.

30.2 OLE between all applications within Groupware

30.3 Understanding how we actually use the tools in a distributed mode. We need to understand process.

30.4 Total integration between all applications.

30.5 how to measure the validity and reliability of groupware tools

30.6 Yes, total integration. Right now artifacts of the tools limit us significantly in ways we need not be limited.

30.7 No matter what we do we will be limited in ways we need not be limited. Thank God for unanswered prayers.

The numbering sequence provides capabilities to uniquely identify each comment. There were 30 separate files that started when the electronic brainstorming tool was initiated. Twenty people each started with an *n.I*, where *n* was a unique number between 1 and 30. Each participant made a comment in response to the electronic brainstorming question and submitted their comment to the public area. They then received another numbered sequence of comments at random (another file), and provided a new comment or commented on the other comments in the file that they received. The ellipses indicates that a substantial number of comments are not included in this sample.

Normal electronic brainstorming stages last for half an hour up to just over an hour. It is easy to see how the brainstorming output can reach hundreds of comments and thousands of lines of text in that period of time with a large group of participants.

The goal of this divergent stage of the process was to identify as many information technology problems as possible during the time allotted for electronic brainstorming. The goal of the next stage, the convergent stage, was to synthesize these comments into a list of problems to be solved over the next five years. The sample of text above demonstrates how difficult this synthesis can be. There are two methods predominantly used to perform this synthesis or idea organization. The method selected depends upon the facilitator's expertise in the subject area of the domain of discussion. If the facilitator does not have knowledge of the subject area, or if the facilitator feels that his or her expertise might bias the group's synthesis of the divergent comments, he or she will do nothing but guide the group through the idea organization process. Otherwise, the facilitator might begin synthesis of the comments during the time the group is generating the comments through identification of key words and synonyms that can be used to begin clustering comments along topical subject areas.

In this instance, since the facilitator was a professor of Management Information Systems, he had expertise in the subject area. Also, while the topic was interesting to all participants, it wasn't very controversial. Therefore, the facilitator felt free to identify key words while the group was generating comments. The group was not aware that the facilitator was generating the key words while they were discussing the electronic brainstorming question.

What follows is a discussion of how the convergent process is viewed from the perspective of the facilitator, the group, and the AI Concept Categorizer.

Facilitator. Facilitators typically browse the group's comments as they are made in order to determine the sensitivity of the group to the topic (are they hostile, interested, ambivalent, etc.?). As certain terms appear with increasing frequency, facilitators may begin the comment clustering process by entering terms such as "network" or "social" as possible

key words that address the electronic brainstorming question. In this way, participants might start with a straw-man list of key words that may serve as topics for discussion in idea organization, instead of starting with an empty list. In our example, the facilitator generated a list of 20 key words with synonyms associated with some of the key words. The following table lists “uncategorized” as the first topic. While not exactly a topic, the uncategorized set of comments are those comments which do not contain any of the key words or synonyms listed as topics 2 through 21.

	Topic	# Associated Comments
1.	uncategorized	87
2.	video	14
	Synonyms: projection	
3.	network	17
	Synonyms: bandwidth	
4.	multimedia	8
	Synonyms: hypertext, multi media, multi-media	
5.	group memory	6
	Synonyms: project memory, repository	
6.	voice	3
7.	culture	19
	Synonyms: style	
8.	language	8
9.	standards	8
10.	distributed/distance issues	14
	Synonyms: distance, different place	
11.	facilitation	19
12.	research methodologies, etc.	9
13.	cost	8
	Synonyms: money	
14.	team	6
15.	reward	4
16.	integration	4
17.	social	6
	Synonyms: societal, society, socio	
18.	performance	9
19.	virtual	6
20.	education	15
	Synonyms: train, learn, teach	
21.	human	24
	Synonyms: people, user, individual, interpersonal	

Facilitator-generated key word lists are subjective. In addition to having their own unique biases with respect to the subject areas, facilitators are hampered by the ability to

browse all of the comments that are generated. They might miss reading many comments that scroll off the facilitator screen before the facilitator has a chance to view them. Thus, the key word list will be incomplete as well as biased vis-à-vis the knowledge and understanding of the facilitator to the subject area. The amount of bias is difficult to measure, but will usually be reflected in verbal discussion of list items. In the above example, 87 comments were still left in the “uncategorized” category. This means that 43% of the 201 comments are still uncategorized. The sum of the associated comments attached to the keywords is 207, which amount is due to multiple counts of those comments that contain more than one of the key words on the list.

Group process. From the group perspective, starting with 70% of the comments already categorized is better than starting with 0%. If starting with nothing but the electronic brainstorming output, each participant would read through the complete set of comments, synthesizing comments into information technology problem areas that need to be solved. As problems are identified and synthesized they can be added to a public screen for all of the other group members to see; or, participants may have a local blank list on which they can collect their synthesized ideas and submit them to the public screen as a separate step.

In any event, information technology problem ideas get posted to the public screen and are discussed either as to the meaning of the terms being posted and/or whether the posting is relevant to the objective of identifying information technology ideas. Postings to the public list occur anonymously during a GSS session. Discussion can also occur anonymously, but more frequently is verbal. Because of this large amount of verbal discussion and negotiation with respect to meaning and appropriateness of terms on the list, the facilitator typically serves as the negotiator among the multiple agents of the group. This relationship is reflected in the label of the arrow between the Group and Facilitator agents in Figure 2.

Participants have the opportunity of attaching supporting comments to the ideas or topics that they propose to the public list. Actual comments from the electronic brainstorming output can be attached or new comments may also be added. In practice, participants seldom do so; or, if they do, the comments attached are somewhat incomplete in that not all supportive comments in the electronic brainstorming output are linked to the topics. Thus, far fewer comments in the electronic brainstorming output actually get linked to the items on the list generated by a group. Any subsequent review of the list items will most likely require a complete review of the electronic brainstorming output as well, in order to more fully understand the items on the list. Better use of the synthesized list of topics created by the idea organization process as an index back into the electronic brainstorming output file is needed. The AI Concept Categorizer supports this need.

AI Concept Categorizer. Key wording is one technique that helps the facilitator manage the large quantity of information generated in the divergent GSS stage. The AI Concept Categorizer supports the facilitator in information management by providing a more objective straw-man list of concepts than can be created by the facilitator using the key wording technique. The AI tool performs a complete analysis of all of the output of the electronic brainstorming stage. Nothing is missed. Since the analysis of the text is statistical, there is no bias with respect to domain knowledge or lack thereof that may result

from a key word list generated by a facilitator. The arrow connecting the Facilitator and AI Concept Categorizer agents in Figure 2 reflect this information management relationship.

In addition, the group may benefit from the AI Concept Categorizer through the links that are formed between the topics contained on the concept list and the comments that generated the concepts. The electronic brainstorming output is thus completely indexed by the list topics, and can therefore be browsed to further explore the concept and better phrase the concepts addressing the strategic plan component objective.

The list of topics generated by the AI Concept Categorizer follows:

	Topic	# Associated comments
1.	technology/dealing/cultural differences/	26
2.	collaborative systems/	8
3.	meetings/distributed/human/	27
4.	linear thread meeting/	2
5.	ai/data/amounts/	6
6.	environments/virtual/	10
7.	voice recognition/	2
8.	tools/culture/	18
9.	people/	9
10.	technologies/	8
11.	information/ability/	10
12.	video/desktop/	8
13.	networks/	7
14.	issues/	7
15.	hardware/	6
16.	distributed meetings/	3
17.	training/	5
18.	groups/	5
19.	applications/	5
20.	distributed environment/	2
21.	comments not linked to any topic	77

A noticeable difference between the list generated by the “expert” facilitator and the AI Concept Categorizer is the difference in terms that comprise the clusters of each list item. In the case of the facilitator, the list item contains a key word and associated synonyms of that key word. In the case of the AI tool, the list item contains associated terms which, taken as unit, comprise a concept. The semantics of the concept is captured in the collection of terms in the AI tool rather than in one key word and the associated synonyms.

6. General observations

Initial observations from implementing prototypes of the AI Concept Categorizer are very encouraging. The tool has been run against dozens of EBS files and has proven to be very

robust. Reactions from experienced facilitators have been very enthusiastic. The output of the AI Concept Categorizer was presented to the group participating in the example above. A separate anonymous electronic brainstorming session was held the following day to collect their comments. Overall, the group was extremely impressed with the output produced by the AI tool. Sample comments from experienced GSS participants in the demonstration include:

This seems like a really big winner....This is the first groupsystems tool I've ever seen where everyone had an immediate use for it the first time they saw it.

I think that helping people with information overload is vitally important. This tool seems to be going in that direction.

The tool may actually help neutralize the perceived influence of an outside facilitator. In this sense, the computer is seen as an unbiased input that can be very useful in a number of group situations. A good facilitator will use this (and other) tools to augment their skills and improve their own as well as the group's productivity. A group leader and members within a group can also find this type of tool useful to assist in information categorization.

Very powerful and absolutely necessary to promote the dynamics of group interaction in these meetings. The target area for this effort is the more difficult and less (sic) interesting for the participants.

I am very happy about what I saw yesterday. The quality of the output of the AI categorizer seemed very close to the one made by "experts."

Thus, our initial experiences in using the tool have produced these benefits:

More negotiation time: Prior to using the AI tool, participants spent much of the time browsing serially through a text file. This may include much time spent in connecting comments to concepts that they create. The AI tool produces a list of concepts with relevant comments already attached (linked). Discussion can begin immediately as to what the concepts mean to the group based upon the underlying comments. As mentioned previously, this discussion (negotiation) is an essential aspect of creating a common group list of concepts into which all group members can buy.

Completeness: The tool indexes every comment contained in the text file. Term frequency thresholds are set to reduce computational complexity. The list of concepts typically

generated index 90% of the comments. The additional 10% of the comments are placed in an “uncategorized” category, which is very manageable.

Consistency: Group participants can be very inconsistent in identifying concepts from comments as they serially go through a large text file. The cognitive load of the large text file is too great. Concepts generated will reflect the context of the text surrounding the current comment and not the overall set of comments. The AI tool provides a statistical analysis of the text file that is consistently applied across the complete text file. The degree to which this statistical analysis supports semantic analysis needs further experimental determination.

Surface of sensitive topics: Certain finger-pointing or discussion of sensitive topics may occur during the divergent activity as a result of the anonymous nature of the activity. Participants may be reluctant to bring these topics to the list for fear of causing potential verbal disagreements. For better or worse, the AI tool will post these concepts if they were discussed during the divergent activity. Thus, negotiation of the sensitive topic will result, if only to agree to table the issue for a later time.

7. Issues for future research

Other issues have resulted from experience using the AI Concept Categorizer tool in experiments and demonstrations.

Cognitive mapping. The creation of a list of concepts with the linkages to the comments that support those concepts is the first stage in creating a cognitive map in a group setting as described by Eden and Ackermann (1993). Their “Strategic Options Development and Analysis” (SODA) decision support system and methods focus on the linkages between strategic issues rather than on the issues themselves. The concepts and linkages to the comments serve as the first step in the development of an aggregate cognitive map for a group. After the system creates the initial concepts and linkages to group comments the group can then strengthen or diminish the coefficients of these linkages to create a more accurate aggregate cognitive map.

Performance. Elements of performance include not only the time necessary for the AI Concept Categorizer to formulate categories but also the efficiency with which categories can be modified and ideas re-associated with revised categories. Formulation of initial categories is a function of several adjustable parameters (indexing thresholds, neural network weighting adjustments, file length, etc.). An additional element of performance from an operational perspective is the ability to conveniently and rapidly respond to facilitator/leader and group suggestions.

Satisfaction. Satisfaction with the AI categorization process and product is a key issue. To date, users have been pleased. The process is compatible with the way groups work in

developing a list of categories. The product has been consistently accepted at least as a strong starting point towards a final set of categories. The tool has yet, however, to undergo the type of extended use in operational contexts that will ultimately determine broader acceptance and use and integration into day to day activities.

Face validity. The credence and face validity of the list generated by the AI categorizer and face validity is crucial. The initial reaction of the facilitator and/or group leader as well as the group members is extremely important. An unrecognizable list or strange categories in the list reduce confidence in the process and invoke discussion that can quickly lead the group astray from its primary objective. On the other hand, if a group sees immediately that there are good categories, the list become a catalyst for positive discussion and additions that quickly move the categorization process forward.

Completeness. An interesting measure of completeness is the degree to which ideas generated are captured in the final list of categories. A large number of uncategorized comments is indicative of incomplete categorization and tends to weaken the credibility of the final list.

Stopword list. The AI Concept Categorizer has a “stopword list” of commonly used terms, that should be overlooked in the search for key concepts. It may become necessary to dynamically tailor the stopwords list as a function of discussion content. For example, in one session the term “group memory” was important to a category reflecting project memory ideas, but the word “group” was on the stopwords list. This resulted in a number of ideas left as uncategorized. Modification of the stopwords list rectified this problem.

Accuracy. Accuracy of the AI Concept Categorizer in terms of coverage of a set of ideas is difficult to judge. It is not sufficient to judge accuracy on the basis of algorithm software alone. From a group perspective, accuracy also entails a sense that the list of categories represents the intent and breadth of the collection of ideas generated. This judgement of accuracy is much broader but ultimately more indicative of tool success. As noted, initial results are promising.

Consistency. The AI Concept Categorizer is consistent, as expected, in terms of repeatedly arriving at the same set of categories given a common input file. A stronger challenge to consistency is the tool’s ability to internally balance ideas between categories. It is not uncommon for an idea to have applicability to more than one category. It is important to see within the context of a group of ideas in a category the sense of cohesiveness and applicability that lends confidence to the usefulness of the tool.

Credibility. Credibility as an issue captures the belief of the group that the list of categories generated adequately portrays session content, and, further, that the tool has added something to the process is beyond the possible attainment of an unsupported group. Initial reaction in terms of comparison of AI Concept Categorizer output to both facilitator and group generated output has been extremely promising.

Robustness. Robustness in this sense is the ability to the AI Concept Categorizer to work across a broad range of file sizes and content variations. Results to date are promising.

8. Conclusion

The complexities of strategic planning are increased when supported by GSS. Increased participation is beneficial to a more complete strategic plan, and buy-in through participation should aid in strategic plan implementation. While GSS may decrease the complexities of managing larger numbers of participants, it can increase the cognitive demands of the participants due to information overload. Groups can easily and enthusiastically generate large numbers of unique ideas using GSS. A group of size 16 to 20 will typically generate over 200 ideas with an average length of four lines each in 30 minutes using electronic brainstorming. The spectrum of ideas generated are typically organized into categories, again using computer-based support. A set of 200 ideas would generally cluster into 15 to 20 parsimonious categories. When successful, the categories are reasonably non-overlapping and roughly similar in degree of breadth and depth.

An important element of successful categorization is participant satisfaction. First, the participants must be confident that the categories represent a coverage of the range of topics that can be useful in further discussion and refinement. Second, the participants must recognize an element of their own work during the categorization process to “buy in” to the final product. Finally, the participants must feel that the categorization activity has proceeded in a reasonably timely fashion with minimal false starts, rework, and confusion.

Successful categorization is a negotiation process that typically involves a blending of facilitated direction, coupled with group participation. A facilitator or group leader will assist the group in identifying candidate categories, with the group continuing to generate additional categories and attaching associated ideas. At times, group members work independently, identifying candidate categories and attaching supporting ideas. At other times, they work in concert to condense, shape, refine, and finalize categories. Computer support for the facilitator and/or leader is complemented by computer support for the participants and face-to-face discussion. We have taken the viewpoint that participants are multiple agents tasked with producing content of a sessions goals (in this case a strategic plan). The facilitator is an agent tasked with managing the process of content production.

In this article we have introduced a third agent, the AI Concept Categorizer. We described how this agent can provide support for reducing the cognitive loads placed on the facilitator and group participants, thus allowing the participants to more quickly begin negotiation on conceptual content. This should almost eliminate information overload frustrations normally associated with convergent stages of the strategic planning processes (those stages requiring the most negotiation within the group). Negotiation on content can then begin. The AI Tool can further provide support for negotiation by providing links to those comments that support the negotiation topics.

Further experimentation is necessary to address some of the issues raised by development of the AI Concept Categorizer. However, results so far have demonstrated that this new AI agent can be extremely beneficial to facilitating and negotiating components of a strategic plan.

References

- Ackoff, R. L. (1981). *Creating the Corporate Future: Plan or Be Planned For*. New York: John Wiley & Sons.
- Chen, H. et al. (1994). "Automatic Concept Classification of Text from Electronic Meetings," *Communications of the ACM* 37(10), 56–73.
- Chen, H. et al. (1993). "Generating, Integrating, and Activating Thesauri for Concept-Based Document Retrieval," *IEEE EXPERT, Special Series on Artificial Intelligence in Text-Based Information Systems* 8(2) (April), 25–34.
- Chen, H., and K. Lynch. (1992). "Automatic Construction of Networks of Concepts Characterizing Document Databases," *IEEE Transactions on Systems, Man and Cybernetics* 22(5) (September/October), 885–902.
- Connolly, T., L. Jessup, and J. Valacich. (1990). "Effects of Anonymity and Evaluative Tone on Idea Generation in Computer-Mediated Groups," *Management Science* 36(6), 689–703.
- Crouch, C. (1990). "An Approach to the Automatic Construction of Global Thesauri," *Information Processing and Management* 26(5), 629–640.
- Dalton, J., and A. Deshmane. (1991). "Artificial Neural Networks," *IEEE Potentials* 10(2) (April), 33–36.
- Dennis, A. et al. (1988). "Information Technology to Support Electronic Meetings," *MIS Quarterly* 12(4) (December), 591–624.
- Dennis, A., J. Nunamaker, and D. Vogel. (1991a). "A Comparison of Laboratory and Field Research in the Study of Electronic Meeting Systems," *Journal of MIS* 7(3) (Winter), 107–135.
- Dennis, A., J. Valacich, and J. Nunamaker. (1991b). "Group, Sub-Group and Nominal Group Idea Generation in an Electronic Meeting Environment," *Proceedings of the 24th Hawaii International Conference on System Sciences* (January).
- Eden, C., and Ackermann, F. (1993). "Evaluating Strategy—Its Role within the Context of Strategic Control," *Journal of the Operational Research Society* 44(9) (September), 853–865.
- Everitt, B. (1980). *Cluster Analysis*, 2nd ed. London: Heinemann Educational Books.
- Jessup, L., and D. Tansik. (1991). "Decision Making in an Automated Environment: The Effects of Anonymity and Proximity on Group Process and Outcome with a Group Decision Support System," *Decision Sciences* 22(1), 266–279.
- Jones, K. S. (1991). "Notes and References on Early Automatic Classification Work," *ACM SIGIR (ACM Special Interest Group on Information Retrieval) Forum* 25(1) (Spring), 10–17.
- Jones, K. S. (1971). *Automatic Keyword Classification for Information Retrieval*. London: Butterworths.
- Grohowski, R. et al. (1990). "Implementing Electronic Meeting Systems at IBM: Lessons Learned and Success Factors," *MIS Quarterly* (December), 368–383.
- Knight, K. (1990). "Connectionist Ideas and Algorithms," *Communications of the ACM* 33(11) (November), 59–74.
- Lewis, C. Michael, and K. Sycara. (1993). "Reaching Informed Agreement in Multispecialist Cooperation," *Group Decision and Negotiation* (September), 279–300.
- Lippmann, R. (1987). "An Introduction to Computing with Neural Networks," *IEEE Acoustics Speech and Signal Processing Magazine* 4(2) (April), 4–22.
- Martz, B., D. Vogel, and J. Nunamaker. (1992). "Electronic Meeting Systems: Results from the Field," *Decision Support Systems* (April).
- Michalski, R., and R. Stepp. (1983). "Learning from Observation: Conceptual Clustering." In R. Michalski, R. Carbonell, and J. Mitchell (eds.), *Machine Learning, An Artificial Intelligence Approach*. Palo Alto, CA: Tioga Publishing Company, 331–363.

- Morrison, J., M. Morrison, and D. Vogel. (1992). "Software to Support Business Teams," *Group Decision and Negotiation* (September), 91–115.
- Nunamaker, J. et al. (1991a). "Information Technology for Negotiating Groups: Generating Options for Mutual Gain," *Management Science* 37(10) (October), 1325–1346.
- Nunamaker, J. et al. (1991b). "Electronic Meeting Systems to Support Group Work," *Communications of the ACM* 34(7) (July), 40–61.
- Poole, M. S., and M. Jackson. (1993). "Communication Theory and Group Support Systems." In *Group Support Systems: New Perspectives*. Jessup and Valacich (eds.), New York: Macmillan, pp. 281–294.
- Post, B. (1992). "Building the Business Case for Group Support Technology," *Proceedings of the 25th Hawaii International Conference on System Sciences* (January).
- Salton, G. (1989). *Automatic Text Processing*. Reading, MA: Addison-Wesley.
- Salton, G. (1978). "Generation and Search of Clustered Files," *ACM Transactions on Database Systems* 3(4) (December), 321–346.
- Simpson, P. (1990). *Artificial Neural Systems: Foundations, Paradigms, Applications, and Implementations*. New York: McGraw-Hill.
- Tank, D., and J. Hopfield. (1987). "Collective Computation in Neuronlike Circuits," *Scientific American* 257(6) (December), 104–114.
- Valacich, J., A. Dennis, and J. Nunamaker. (1992). "Group Size and Anonymity Effects on Computer-Mediated Idea Generation," *Small Group Research* 23(1), 49–73.
- Vogel, D., and J. Nunamaker. (1990). "Group Decision Support System Impact: Multi-Methodological Exploration," *Information and Management* (Winter), 15–28.
- Vogel, D. et al. (1989). "Electronic Meeting System Experience at IBM," *Journal of MIS* 6(3) (Winter 1989–1990), 25–43.