

Extracting team mental models through textual analysis

KATHLEEN M. CARLEY

Department of Social and Decision Sciences, Carnegie Mellon University, U.S.A.

Summary

An approach, called map analysis, for extracting, analyzing and combining representations of individual's mental models as cognitive maps is presented. This textual analysis technique allows the researcher to extract cognitive maps, locate similarities across maps, and combine maps to generate a team map. Using map analysis the researcher can address questions about the nature of team mental models and the extent to which sharing is necessary for effective teamwork. This technique is illustrated using data drawn from a study of software engineering teams. The impact of critical coding choices on the resultant findings is examined. It is shown that various coding choices have systematic effects on the complexity of the coded maps and their similarity. Consequently, a thorough analysis requires analyzing the data several times under different coding choices. For example, re-analysis under different coding scenarios revealed that although members of successful teams tend to have more elaborate, more widely shared maps than members of non-successful teams, this difference is significant only when the data is unfiltered. Thus a better interpretation of this result is that all teams have comparable models, but successful teams are able to describe their models in more ways than are non-successful teams. © 1997 John Wiley & Sons, Ltd.

J. Organiz. Behav. 18: 533-558 (1997)

No. of Figures: 3 No. of Tables: 4 No. of References: 85

Introduction

Recently there has been an upsurge of interest in individual and team mental models (Reger and Huff, 1993; Johnson-Laird, 1983; Klimoski and Mohammed, 1994; Eden, Jones and Sims, 1979; Carley, 1986a; Roberts, 1989; Weick and Roberts, 1993; Walsh, 1995). On the one hand, differences in individual mental models are potentially valuable for organizational performance enabling the organization to learn from different individual's experiences (Knorr-Cetina, 1981; Latour and Woolgar, 1979). On the other hand, common team mental models are seen as critical for team learning and team performance (Hutchins, 1990, 1991a,b; Fiol, 1994). The relation of individual mental models to team mental models, and the value of team models to team and organizational performance is the subject of much debate. Teams are being used in organizations with increasing

This work was supported in part by the NSF under grant IRI-9216760. Many thanks to Sara Kiesler and Doug Wholey for their involvement in the larger project and for their comments on an earlier version of this paper. I would also like to thank the reviewers for their extensive and very helpful comments. The comments of one reviewer in particular were incredibly valuable not just for this paper, but for future work in this area.

frequency. However, the conditions for team success are not well understood. Thus the cognitive processes underlying team behavior are in need of investigation. Indeed, a variety of important organizational issues can be cast in terms of team mental models. For example, to what extent must individuals share their mental models if they are to operate effectively as a team? Do successful teams have different team mental models from non-successful teams? Or, are there certain types of information that more commonly appear in team models? While there is general agreement that it would be useful and important to examine team mental models there is little understanding as to how to uncover, derive, code, or analyze such models. We must address this methodological question before we can even begin to address these important theoretical questions.

Recent advances in textual analysis provide a technique that allows the researcher, given a set of texts¹, to extract, code and analyze individual and team mental models as cognitive maps² (Carley and Palmquist, 1992; Roberts, 1989, 1997; Eden *et al.*, 1979; Huff, 1990; Walsh, 1995). Previous work using cognitive mapping typically required the researcher to code texts using, at best, a semi-automated procedure. Sets of maps, extracted from texts, could then be compared. In this paper, an automated approach to deriving team mental models is presented and illustrated. This approach builds on, and extends, the semi-automated procedures developed by Carley (1993) for coding texts as maps. This paper contains two methodological contributions. First, an automated approach for coding texts as cognitive maps is presented and the impact of key coding choices on the resultant maps is demonstrated. This automated approach is embodied in the program—*automap*. Second, a procedure for combining these coded individual cognitive maps to derive the team map is described. This procedure makes use of standard UNIX utilities. Then, these tools for coding texts as cognitive maps are illustrated using a sample of data from an extensive study of software engineering teams collected by Kiesler, Wholey and Carley (1994). In order to illustrate the approach for, and the potential power of, deriving and comparing team maps, the issue of whether there are systematic differences in successful and non-successful teams is cursorily explored. The results should be taken as indicative of the types of analyses made possible by this approach to team mental models, illustrative of the resultant strengths and weaknesses of this approach, and not as providing conclusions about the determinants of team success. More importantly, from a methodological standpoint, in doing this illustration the relative impact of different coding choices on the robustness of the substantive conclusions are explored.

¹ The term text is used in its most generic form to refer to any symbolic data that is concept based either written or verbal and that has both a syntactic and semantic structure. Thus, answers to open-ended interview questions, speeches, television shows, dialogues, newspaper articles, books, open ended survey questions, correspondences, all are considered as text. Clearly there are differences between oral and written texts. The proposed method can be used to study some of these differences.

² The term cognitive map has been widely used in a number of literatures. Typically it refers to some form of network-based representation of an individual's cognitions. These cognitions may be about other companies or people and the relations between them (Daniels, Johnson and de Chernatony, 1994), or about ideas and the relations between them (Eden *et al.*, 1979; Roberts, 1997), or both (Carley and Palmquist, 1992). The term map analysis is used herein to differentiate the proposed approach from the myriad of other techniques for constructing and analyzing cognitive maps. Two such commonly used techniques are the repertory grid technique (Kelly, 1955; Bannister and Mair, 1968) and the visual card sorting technique (Daniels *et al.*, 1994). A repertory grid consists of a matrix where rows are 'constructs' and columns are 'elements'. Constructs are dichotomous scales (e.g. qualitative/quantitative). Elements are exemplary cases which respondents rate, say, from 1 to 5, in terms of each construct. The resulting matrix can then be used to calculate distance measures, conceptual clusters, etc. Repertory grids can be compared across individuals particularly when a predefined list of concepts are used. Visual card sorting involves first locating a set of ideas (or names of people or companies) (possibly from the respondents), then writing one idea per card, and finally as part of a structured interview asking the respondents to sort the cards and position them as to what is closest to what. In contrast to these two techniques, the proposed technique can be used to extract a cognitive map from any one type of text (thus it can be used with non-invasive and non-reactive data collection techniques), can require less data collection time (and, as compared to other textually-based data analysis techniques requires less time in analysis), and is flexible enough to be used on a variety of definitions of team mental model.

What is a mental model?

Mental models have been characterized as internal symbolic representations of the world or aspects of the world (Johnson-Laird, 1983). An individual's mental model of some topic includes the individual's definitions, procedures, examples, and so on. Individuals use these mental models to negotiate their lives, determine which actions to take, and construct the social world (Mead, 1962; Blumer, 1969; Goffman, 1974; Stryker, 1980; Walsh, 1995). Individuals have many mental models, utilize different mental models at different times, and adapt or alter their mental models over time.

At the individual level there is debate over the degree to which individuals are aware of or have preconstructed their own mental models. On the one hand, some consider these models inherently unobservable as they are 'in the mind' of the individual and employ tacit and unarticulated knowledge. From this perspective the researcher can only code and analyze an elicited representation of the mental model—a cognitive map. For others, the mental model is an emergent structure and only comes into being as the individual articulates it and is therefore observable (see discussions of emergent cognition, Sperry, 1995; Hergenbahn, 1994; Forman, Minick and Stone, 1993). Further, this articulation is expected to be incomplete for a variety of reasons including attention and awareness. From this perspective the mental model (as emergent structure) and the cognitive map (barring coding errors) are in practice indistinguishable.

For teams, the team mental models represent 'shared' or 'social' knowledge. However, at the team level, the nature of team mental models is in dispute. There are three major areas of contention: uniformity of sharing—whether knowledge must be uniformly shared by team members (the relationship between individual and team knowledge); degree of sharing—how widely the knowledge must be shared; and awareness of sharing—whether or not the team members must be aware that the team model is shared.

For some, information is only in the team model if it is uniformly shared. For example, Klimoski and Mohammed (1994) discuss the team model as a set of knowledge that all know. In contrast, for Polanyi (1958, especially, pp. 216–219 and 264–266) social knowledge requires a transitivity of appraisal across a continuous network of individuals. What this means is that each 'piece' of social knowledge is commonly, but not necessarily uniformly, shared. Similarly, both Durkheim (1938) and Fauconnier (1985) talk about the social existence of a collective mind and suggest this entity has a life of its own. For example, Durkheim (1938) talked about the group mind or the collective mental state of the group members as a social entity whose impact and constitution extended beyond that of a single individual and could not be completely mapped onto a single individual. These notions of social knowledge and collective mind gain precision when we define them as the team mental model. From this perspective, team mental models represent only a lossy³ consensus to a set of information as some of the information that is known to the society as a whole is lost over time or across groups as the set of individuals over which consensus is measured changes (Carley, 1997). Since this model is created through a lossy integration process it will be impacted by each individual, yet have a life-time independent of any

³ Lossy consensus is a consensus that is generated by a lossy integration process. A lossy integration process is one in which information is lost as, for example, when taking a moving average. A lossy consensus is one in which what is shared changes gradually over time through a moving average process. For team maps, the basic idea is that each individual's mental model can be represented as a cognitive map. Teams maps can be represented as a composite map by combining (somehow) the cognitive maps of all team members at that point in time. Since individuals may enter and leave teams, and since individuals can learn and so evolve their mental models or fail to articulate portions of their model the individual cognitive maps change. Consequently, the team map is lossy.

particular individual, and taken by itself might suggest a logic of action distinct from what any individual alone might engage in.

If the perspective taken is that the team mental model does not require a uniformity of sharing then the issue arises as to what degree of sharing is needed. One perspective comes from thinking of the team map as culture—those ideas, beliefs, norms, ways of acting, etc. that are commonly though not necessarily universally shared (Martin, 1992; see also Barley, 1983). Martin considers information to be part of this model when 50 per cent or more of the group shares it. Similar arguments are put forward by Krackhardt (1987), Carley (1984, 1986a) and Romney, Weller and Batchelder (1986) in discussions of cultural consensus. The higher the degree of sharing imposed the smaller the resultant team model (Carley, 1986a).

A final issue is whether or not the team members are aware that they share information. For Polanyi (1958, especially, pp. 216–219 and 264–266) for knowledge to be social awareness of its commonality is not needed; rather, all that is needed was 'tacit consensus' which results from the articulation of 'common experience'. The team model exists whether or not any single individual is aware of it, in its entirety. Weick and Roberts (1993) suggest that it is only necessary for the team members to be 'heedful' in their interactions with each other for a team model to emerge. In contrast to these perspectives Klimoski and Mohammed (1994) suggest (p. 422) that a team mental model is an emergent group phenomenon and that since these team mental models facilitate group coordination and the allocation of resources 'some level of awareness is necessary'. They argue that explicit and not simply tacit agreement is needed in order for a piece of information to be part of the team mental model.

The techniques that will now be described are flexible enough to be used regardless of the researcher's perspectives on the foregoing issues (see also Carley and Palmquist, 1992; Carley, 1993; Carley and Kaufer, 1993). To be sure these different perspectives would lead to differences in what data is collected, how it is collected, and what interpretation is placed on the resultant coded maps. However, these perspectives would not affect the general coding technique. In this paper, the focus is on the methodology for creating and analyzing maps and how coding choices influence the robustness of the results.

Coding individual mental models

Texts can be thought of as containing a portion of the author's mental model at the time the text was created (Kaufer and Carley, 1993). Textual analysis techniques can be used to extract and analyze these samples of the author's mental model. A wide number of textual analysis techniques have been used in many different fields⁴ to extract a representation of a mental model from a text (Huff, 1990; Walsh, 1995; Roberts, 1997). The most commonly used textual analysis techniques⁵ used in this way are: content analysis (Namenwirth and Weber, 1987; Stone and

⁴ Indeed advances in textual analysis techniques are occurring in a large number of fields including: political science, sociology, management science, marketing, organizations, communication, rhetoric, and artificial intelligence. Many, although certainly not all, of these techniques are designed to help extract some type of representation of the individual's mental model.

⁵ Less commonly used are the semantic-based techniques such as semantic grammars (Franzosi, 1990a,b) and linguistic content analysis (Roberts, 1989). These techniques have the advantage that the results can be analyzed statistically. However, unlike the techniques mentioned above their aim is to understand the story within the text and not the underlying mental model.

Cambridge Computer Associates, 1968a; Stone, Dunphy, Smith and Ogilvie, 1968b; Ogilvie, Stone and Kelly, 1982), procedural mapping techniques (e.g. semantic-planning nets (VanLehn and Brown, 1980), procedural task analysis (VanLehn and Garlick, 1987)), protocol analysis (Ericsson and Simon, 1984; Newell and Simon, 1972), and various cognitive mapping techniques (Axelrod, 1976; Eden *et al.*, 1979; Shavelson, 1972; Reitman and Rueter, 1980; Carley and Palmquist, 1992; Carley, 1993)⁶. At one level these three techniques have the same goal—extract a representation of the mental model from the text. However, since the basic conception of what is a mental model is different for each of these approaches, the techniques are actually quite different and what information they extract is different.

Researchers using content analysis tend to assume that the author's mental model can be adequately represented by keeping track of the presence of, absence of, and frequency with which certain concepts are used in the text. It is assumed that the meaning of concepts is not affected by context. Consequently, this technique focuses on counting what words or general concepts are present in a text. In contrast, researchers employing procedural mapping typically assume that a mental model is a set of implicit or explicit procedures used by a speaker or author to perform a given task and the order of application. Consequently, this technique focuses on the actions of the author, the order of processing, and the logic of the reasoning. Finally, cognitive mapping researchers typically assume that the mental model is some set of 'cognitive structures' which either exists in memory or emerges during interactions and may include both declarative and procedural information, and the meaning and use of that information depends on the way in which it is interconnected. For example, Shavelson (1972, pp. 226–227) defines the cognitive structure as 'a hypothetical construct referring to the organization (relationships) of concepts in memory'.

Researchers use cognitive mapping techniques in which both concepts and the relationships between them are coded because they are concerned with the extraction of meaning, and they assume that context affects meaning. Researchers use cognitive mapping techniques to extract both concepts and relationships so that they can determine whether the text's authors are just using the same words or are actually exhibiting shared meaning. Cognitive mapping, in a sense, combines elements of both content analysis and procedural mapping. In coding texts, the researcher using cognitive mapping technique will, like the content analyst, determine what concepts are present; but unlike the content analyst will also extract the relationships between those concepts. Further, in a cognitive map, the procedures can be represented as sets of related concepts; however, depending on the researchers particular interest, only declarative and not procedural data might be coded.

In this paper, an automated approach for extracting the cognitive map from a text is presented. This approach uses the program *automap*. The cognitive mapping approach used in this paper is based on Carley's (1986a,b, 1988) map analysis⁷ techniques for coding mental

⁶ The goals and techniques associated with these alternative coding schemes, and the differences in them are described in more detail in Carley and Palmquist (1992). Many techniques for coding texts do not appropriate for handling large numbers of texts. In particular, many of the AI based techniques provide detailed codings of only a few texts (see for example, schemes (Anderson and Bower, 1973; Bobrow and Norman, 1976), schemata (Rumelhart and Ortney, 1976), structured frames (Charniak, 1972; Minsky, 1975), and scripts (Schank and Abelson, 1977) and do not admit statistical analysis of the results. Also many of the most interesting ideas for thinking about texts as maps, such as those by Sowa (1984) on conceptual structures, have not been operationalized into coding tools.

⁷ The techniques used herein build on and extend the MECA tools for map analysis reviewed in Weitzman and Miles (1995). For some readers the term map analysis might be seen as misleading or not the best choice of terms as the tool described is one of a class of tools for doing cognitive mapping. A more precise name might be automated cognitive mapping using proximally linked concepts. Within the area of textual analysis there are a large number of terms for similar, and sometimes even identical constructs (see for example the comment by Walsh, 1995). This is natural as the

models as cognitive maps. Carley previously provided a semi-automated approach to coding texts as networks of concepts by first generalizing the concepts that appear in the text and then coding the set of relations among pairs of concepts. In keeping with the work in artificial intelligence she referred to the coded text as a map and each pair of concepts linked by a relation as a statement. In this paper a procedure that automates the extraction of the map is presented. Then, a second procedure for deriving the team map is presented and illustrated. Only a brief review of the pertinent aspects of the coding scheme and key definitions will be provided. Carley's approach represents a blending, and in some sense a generalization of and operationalization of, decision networks (Axelrod, 1976; Eden *et al.*, 1979), semantic networks (Schank and Colby, 1973), augmented transition networks (Winston, 1977; Clark and Clark, 1977; Collins and Loftus, 1975; Wyer and Carlston, 1979), and conceptual networks (Sowa, 1984). The limited cognitive underpinnings of this approach at the individual level and its relationship to alternate approaches to coding texts from a mental model perspective can be found in Carley (1988, 1993) and Carley and Palmquist (1992).

To orient the reader, however, this approach to coding texts as maps is graphically illustrated in Figure 1. Figure 1 contains a short text and the resultant coded maps (one when no concepts are generalized or deleted and one when concepts are generalized and deleted) that graphically illustrates the key terms in this coding scheme—concept, relationship, statement, and map. Following is a brief definition of these key terms as they are used in the automated approach to coding texts as maps. Critical coding issues are discussed. Then, the automated approach for coding texts as maps is described.

Concepts

A concept is a single ideational category. A concept can be a single word such as 'organization' or 'process', a composite word such as 'information system', or a more complex phrase such as 'democratic and equal participation'. In Figure 1 each concept is denoted by a single word.

Relationships

A relationship is the tie that links two concepts together. These relationships may be conceptual such as in 'is-a' links or may simply denote that the concepts are somehow cognitively related in the mind of the text's author or are proximally located in the text (Axelrod, 1976; Minsky, 1975; Schank and Abelson, 1977; Winston, 1977; Clark and Clark, 1977; Collins and Loftus, 1975; Wyer and Carlston, 1979). For additional details on how to code relationships see Carley and Palmquist (1992). In Figure 1, a relationship is denoted by a line between two concepts. In Figure 1, the only relationship shown is proximity; i.e. concepts that are immediately contiguous in the text have a relationship between them.

textual analysis tools are being concurrently developed by many researchers in a large number of scientific disciplines. Until some cross-disciplinary agreement is reached on tools and terms, differences of opinions on what to call various methods is to be expected. All this being said, in this paper the term map analysis is retained not only because it has appeared previously in the literature but also so that the reader can trace the proposed method to its intellectual and methodological roots in sociology, cognitive science and artificial intelligence.

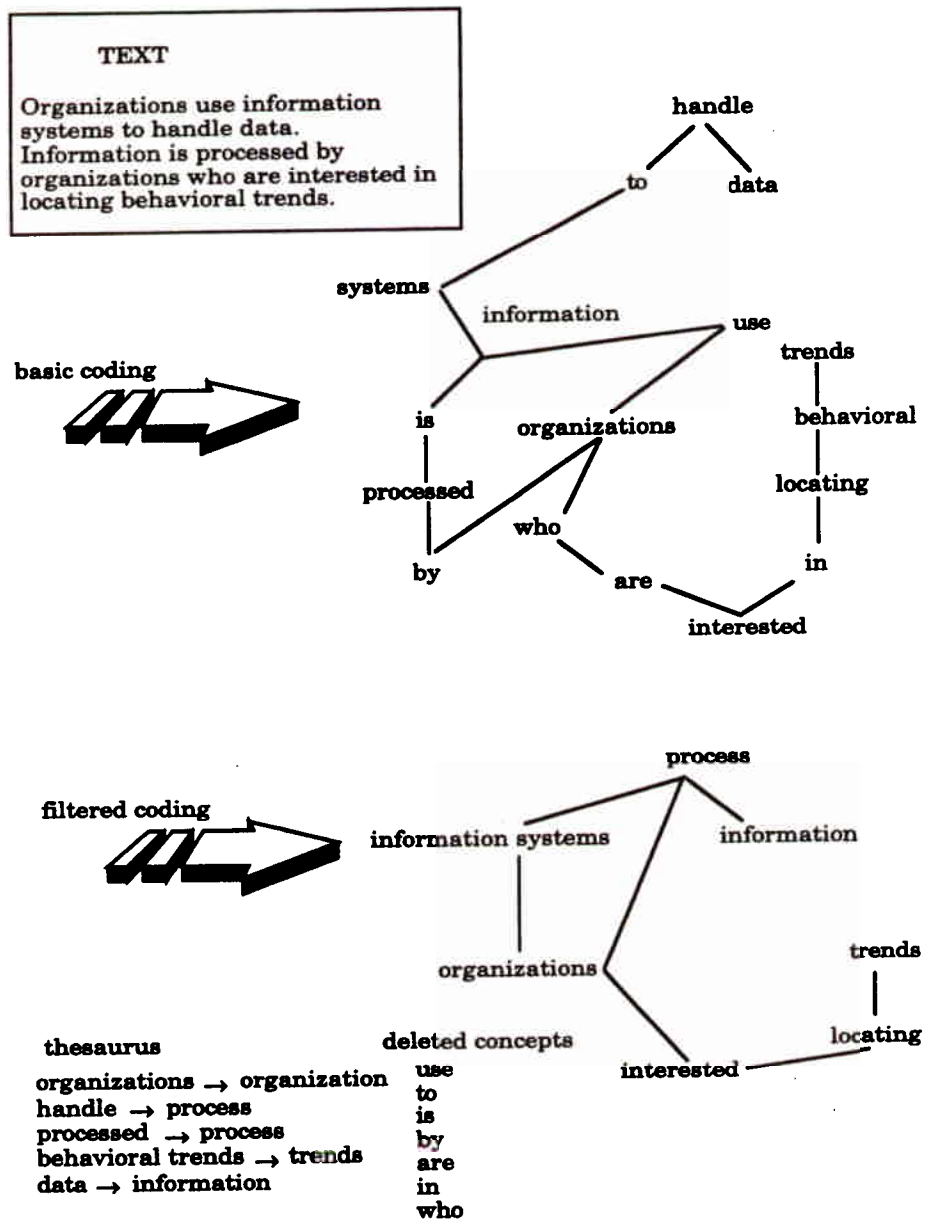


Figure 1. Illustrative text and maps

Statements

A statement is two concepts and the relationship between them. Illustrative statements are 'organizations have cultures' or 'information is processed'. Both statements can be coded as a relationship between two key concepts 'organizations–cultures' and 'information–process'. In Figure 1 each pair of concepts connected by a line is a statement.

Maps

A map is a network of concepts formed from statements. Two statements are linked if they share one concept. For instance, the two statements 'information-process' and 'process-organizations' share the concept 'process' thus resulting in the network 'information-process-organizations'. A map that is coded from a text generated by an individual can be thought of as a representation of that individual's mental model. A map created by intersecting two individual's maps can be thought of as a team map—i.e. a representation of the shared or team mental model. In Figure 1, two maps are shown—one for the unfiltered and one for the filtered text.

Critical coding issues

In order to code a text the researcher must make a large number of coding choices (Carley, 1993). These choices can dramatically alter the resultant analysis. Rather than provide an exhaustive summary, I want to focus on two issues that are particularly critical when automated facilities for coding texts as maps are used. These issues center around the ideas of filtering and windowing.

Filtering involves making a series of critical decisions about which part of the text to code and what words to use in the coding. At a gross level, researchers filter texts by deciding which portion to code (e.g. the whole text or just certain paragraphs). At a more detailed level, filtering involves selecting to code only specific concepts. That the researcher must determine whether all concepts will be used or whether some will be deleted, is critical. The researcher also needs to determine whether the exact words in the text will be used or whether phrases and words will be generalized using some type of thesaurus. Generalization can be used to move the coded text beyond explicitly articulated ideas to implied or tacit ideas. Generalization can also be used to avoid concept misclassification due to peculiar wording on the part of individuals. Deletion and generalization can facilitate the comparison of texts. For example, in Figure 1, the text is shown coded first using each word as a separate concept (basic coding—top) and secondly after some words have been deleted or generalized (filtered coding—bottom). Note, in the filtered coding composite words such as 'information system' are used. In this case, the filtered coding clarifies the relationships between key concepts and admits greater comparison between maps.

There is an important tradeoff here. Deletion and generalization can lay bare the relationships between key concepts and help the coder to extract meaning. Such filtering, can increase the possibility of locating similarities between maps. However, it does so at the cost of individual differences. It is possible to over-generalize and over-delete to the point where meaningful individual differences are lost. Currently, there is no mathematical criteria for determining the appropriate level of filtering. Such a decision must be made by the researcher taking into account the data and the research question being addressed. Possibly some type of inter-coder reliability could be used. Two or more coders could analyze a subset of texts and independently suggest a level of filtering. The extent to which they agree could substitute for a mathematical criterion. In this case, the subjective judgment being made is about 'when is critical information about the topic of concern not being lost due to filtering?' In contrast, typical inter-coder reliability tests look at the question 'does this word or phrase mean the same thing as the word or phrase in the researcher's list of key concepts?' To be sure, with respect to generalization, inter-coder reliability could be done in the classical fashion on a concept-by-concept basis looking at whether or not there is agreement on whether specific concepts should be generalized to other concepts. However, this would not address the filtering due to concept deletion.

Maps

A map is a network of concepts formed from statements. Two statements are linked if they share one concept. For instance, the two statements 'information-process' and 'process-organizations' share the concept 'process' thus resulting in the network 'information-process-organizations'. A map that is coded from a text generated by an individual can be thought of as a representation of that individual's mental model. A map created by intersecting two individual's maps can be thought of as a team map—i.e. a representation of the shared or team mental model. In Figure 1, two maps are shown—one for the unfiltered and one for the filtered text.

Critical coding issues

In order to code a text the researcher must make a large number of coding choices (Carley, 1993). These choices can dramatically alter the resultant analysis. Rather than provide an exhaustive summary, I want to focus on two issues that are particularly critical when automated facilities for coding texts as maps are used. These issues center around the ideas of filtering and windowing.

Filtering involves making a series of critical decisions about which part of the text to code and what words to use in the coding. At a gross level, researchers filter texts by deciding which portion to code (e.g. the whole text or just certain paragraphs). At a more detailed level, filtering involves selecting to code only specific concepts. That the researcher must determine whether all concepts will be used or whether some will be deleted, is critical. The researcher also needs to determine whether the exact words in the text will be used or whether phrases and words will be generalized using some type of thesaurus. Generalization can be used to move the coded text beyond explicitly articulated ideas to implied or tacit ideas. Generalization can also be used to avoid concept misclassification due to peculiar wording on the part of individuals. Deletion and generalization can facilitate the comparison of texts. For example, in Figure 1, the text is shown coded first using each word as a separate concept (basic coding—top) and secondly after some words have been deleted or generalized (filtered coding—bottom). Note, in the filtered coding composite words such as 'information system' are used. In this case, the filtered coding clarifies the relationships between key concepts and admits greater comparison between maps.

There is an important tradeoff here. Deletion and generalization can lay bare the relationships between key concepts and help the coder to extract meaning. Such filtering, can increase the possibility of locating similarities between maps. However, it does so at the cost of individual differences. It is possible to over-generalize and over-delete to the point where meaningful individual differences are lost. Currently, there is no mathematical criteria for determining the appropriate level of filtering. Such a decision must be made by the researcher taking into account the data and the research question being addressed. Possibly some type of inter-coder reliability could be used. Two or more coders could analyze a subset of texts and independently suggest a level of filtering. The extent to which they agree could substitute for a mathematical criterion. In this case, the subjective judgment being made is about 'when is critical information about the topic of concern not being lost due to filtering?' In contrast, typical inter-coder reliability tests look at the question 'does this word or phrase mean the same thing as the word or phrase in the researcher's list of key concepts?' To be sure, with respect to generalization, inter-coder reliability could be done in the classical fashion on a concept-by-concept basis looking at whether or not there is agreement on whether specific concepts should be generalized to other concepts. However, this would not address the filtering due to concept deletion.

Windowing is a process of locating statements by placing relationships between pairs of concepts that appear within a window. A window is a set of contiguous concepts. For example, the entire last sentence could be contained in a window of size 8. But when exactly are two concepts related? Various authors have thought of relationships in temporal (Schank and Abelson, 1977), conceptual (Axelrod, 1976; Minsky, 1975), or proximal terms (Danowski, 1982, 1988). There is little agreement over the types or categories of relations. However, in all cases, the argument is made that the pattern of linkages is important. And, in all cases, the only concepts that are connected by a relationship are those that are physically proximal (often in the same sentence or the following sentence) or are inferred through a generalization process from concepts that are physically proximal. The researcher needs to determine which concepts will have relationships placed between them and how proximally distant concepts can be from each other and still have a relationship between them. At one level, this is simply a matter of determining how close is close enough when placing a link between concepts. At another level, this is a matter of determining the appropriate interplay between filtering and distance. For example, with less filtering a larger window may be needed to access linkages among critical nouns and verbs. Whereas, with more filtering since irrelevant concepts (such as articles) have already been deleted, a smaller window may be sufficient. What the appropriate tradeoff is between filtering and windowing is a point for additional research.

Automap—automating the coding of individual cognitive maps

In this paper, individual's cognitive maps are extracted from texts using a program called automap⁸. This approach combines Carley's semi-automated approach to coding texts as cognitive maps with Danowski's (1982, 1988) completely automated approach for proximity analysis. Carley (1986a, 1988) presented a semi-automated approach to coding text in which the researcher coded a text by first determining which concepts to overlook and which to generalize and then coded the text statement by statement and chose which pairs of concepts to put relationships between. The advantage of Carley's semi-automated approach is that it stays close to the data and lets the researcher make informed or expert choices about what statements exist, how to generalize concepts, which concepts to overlook, and which multi-word phrases to treat as a single concept. This approach does not blindly place relationships by position. This is particularly important in complex or compound sentences. The disadvantage of this approach is that it is slow and cumbersome when the researcher wishes to code a vast number of texts (although it is better than a totally non-automated approach). For this type of coding inter-rater reliability is typically higher for locating the same concepts than for locating the same statements (Carley, 1988).

Danowski (1982, 1988) presented a completely automated approach to coding text in which the computer places a link between each pair of words within a window, i.e. within a certain number of words of each other. In this case, the texts were unfiltered and all words were treated as distinct concepts. The advantage of Danowski's automated approach is that it is fast and it eliminates the issue of inter-rater reliability with respect to locating the same statements. The disadvantage is that it blindly puts links between all words within a given window without consideration for

⁸ Automap is a program written in C, that runs under the UNIX operating system, and is available as is from the author. It takes as input the original text (ASCII format), a list of concepts to delete, and a thesaurus file showing which concepts to translate into which others. Automap generates two output files, the filtered text and the coded map. All files are ASCII files.

alternate word meanings. Further, since the texts are unfiltered many of the statements reflect grammatical style; e.g. statements linking 'and' and 'the' are quite common. Further, Danowski's approach relies on the assumption that words with similar meanings will exhibit similar patterns of connectivity. This assumption may be unjustified, particularly when one is trying to derive the team map. However, this is a point that needs further research.

As noted, automap combines these two approaches. The proposed approach requires a two-step process: (1) filtering and (2) windowing. First, the user provides a list of concepts to delete and a thesaurus denoting which concepts get changed into which other concepts. Automap filters the texts by applying a preprocessor to the text that deletes those concepts in the researcher's delete list and then uses the researcher's thesaurus to translate words in the text. The resulting text is a filtered version of the original text provided to automap. Second, automap uses a windowing process to automatically generate a map from this filtered text. Specifically, automap codes the text as a map by putting a relationship between each pair of concepts that occur within a window. The size of the window is defined by the researcher. The maps in Figure 1 were generated using a window of size 2.⁹ The resulting map is a file with one statement per line. This map can be equivalently thought of as a symmetric binary graph where each concept represents a node and each relationship represents a bi-directional tie between the two concepts.

The proposed combined approach, as embodied in automap, is a relatively fast approach to coding texts. Through the filtering process it lets the researcher make informed or expert choices about how to generalize concepts, which concepts to overlook, and which multi-word phrases to treat as a single concept. Through the windowing process it eliminates the issue of inter-rater reliability with respect to locating the same statement. The combination of filtering and windowing reduces the likelihood that words with alternate meanings are misconnected. Also, judicious deletion as part of the filtering process eliminates the likelihood of statements reflecting grammatical conventions, such as 'that-is' and 'and-the', being coded. The disadvantage of this approach is that it may overly link concepts which are not semantically related simply because they are contiguous. Moreover, and possibly of greater importance, relationships between concepts that are distant from each other due to grammatical phrasing but are still meant to relate to each other will not be picked up. It appears possible to fine tune this approach and largely overcome these potential difficulties by judicious deletion of uninteresting or irrelevant words, preprocessing the text using a thesaurus that contains generalizations, altering the window size, and converting pronouns to specific nouns. Initial investigations look promising; however, this is an issue which needs further research¹⁰.

⁹ The window floats across the text so that any pair of concepts within 'x' of each other are linked into a statement. Imagine a window of length 4. Consider the sentence 'Information can be easily retrieved, stored, and altered'. This would result in 18 statements assuming each word is a separate concept. These statements are information-can, information-be, information-easily, can-be, can-easily, can-retrieved, be-easily, be-retrieved, be-stored, easily-retrieved, easily-stored, easily-and, retrieved-stored, retrieved-and, retrieved-altered, stored-and, stored-altered, and-altered.

¹⁰ Other than these factors, the main thing that is lost in going from the semi-automated to the more automated approach is coding flexibility. For example, in the semi-automated approach (Carley, 1993; Carley and Palmquist, 1992) meta-classification schemes can be used to categorize concepts and relationships into types, and relationships can be further characterized in terms of direction, sign and strength. Since these alternative coding options were rarely used by people coding texts as cognitive maps they were not implemented in the automated program (although it could be readily extended to encompass these details). Whether this loss in coding flexibility is critical is a point for further research. More importantly, this added coding flexibility greatly increases the complexity of constructing the team cognitive map. The scheme presented above amounts to coding the text as a symmetric binary graph. The statistical and mathematical techniques for analyzing such binary graphs have been adequately developed. A more flexible coding scheme, for example one that allowed concept types or relationship weights, would result in colored or non-binary networks. Further mathematical and statistical research is needed for analyzing multiple weighted or colored networks. Thus, in this paper, the focus is on maps coded in the most simple fashion.

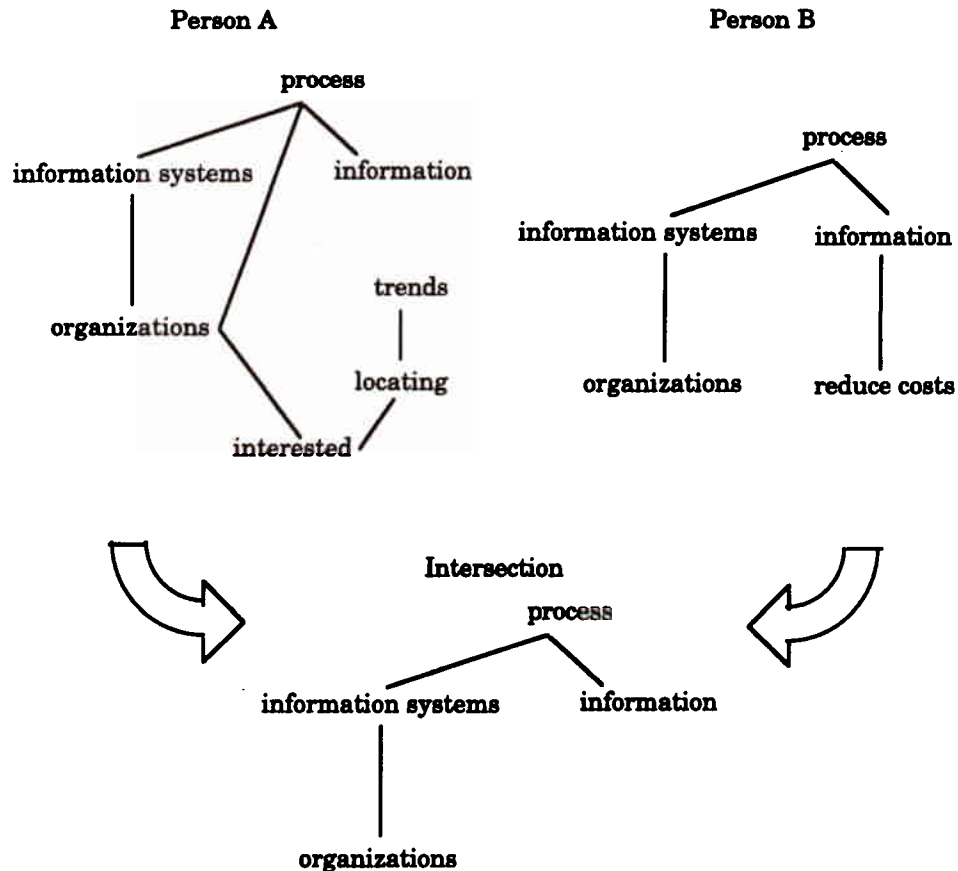


Figure 2. Intersection of two maps

Constructing team cognitive maps

Individual cognitive maps can be compared, or combined to create a team cognitive map. Assume that we have coded two individual's mental models as cognitive maps. The resultant maps are shown in Figure 2. How can we compare and contrast these maps? Each individual cognitive map can be thought of as a binary graph. As such we can compare them simply by counting the number of shared concepts, shared statements, total concepts, total statements, concepts only in that map, statements only in that map. We can combine two maps by creating either a union or intersection file. For example, in Figure 2, the bottom map represents the intersection of the two maps on the top. This intersection, is also a map, and can be thought of as a representation of the shared mental model.

When there are more than two individuals in a team, deriving a representation of the team mental model, a team map, involves intersecting multiple maps. Depending on the level of generalization done when the individual cognitive maps were coded, these team maps may contain only that information that is explicitly articulated by all members, or tacit information

that is implied by the articulated information. Depending on the perspective taken on uniformity of sharing either a union, an intersection, or a lossy intersection rule can be used. Lossy intersection means all those statements that are shared by at least 'x' team members are part of the team map. Here the researcher sets 'x' as some cutoff value¹¹. The desired degree of sharing is reflected in how many individual maps a statement must occur in before it appears in the team map.

Data for illustrative example

Sample

To illustrate these techniques, abstractions of the mental models of all of the students in an information systems project course at a private university were used. There were 41 students. In this course the students were divided into eight small project teams (four to six members) whose goal was to analyze a client's need, and then design and build an information system to meet that need within one semester for a 'pretend' client. Five of the teams are considered to be successful in their project development, and three are not. Team success was determined by a combination of client input and instructor rating. Three times during the term (beginning, middle, and end) the students were asked to answer an extensive questionnaire. On this questionnaire two open-ended questions were asked 'What is an information system?' and 'What leads to information system success and failure?' Students' cognitive maps are coded from the answers to these two questions at the beginning of the term. Thus, the presented maps can be thought of as representations of individuals' initial mental models based on the information they have articulated.

The students are in their junior year and this is their first group experience. Students remained in the same team throughout the term. Effort was made to 'equalize' the teams in terms of intelligence and competency. Specifically, groups did not differ significantly in the grade point average of their members. Students with exceptional programming ability or experience were distributed uniformly across the teams. Students with exceptional background in writing were distributed uniformly across the teams.

Each team in this study had a manager which they as a team selected. Often, this manager was the student whom the students expected to have had the most previous experience in designing and building information systems. Although, in actuality the team managers did not have substantially more experience and certainly were not 'experts' in the subject area. Team managers were not necessarily more intelligent or verbal than other team members (based on grade point average). The role played by the team manager varied by group. In general, the manager tended to coordinate the group, assign subtasks, and call team meetings. Often they spent more hours on the project (based on weekly self-reports of hours spent).

This data was collected as part of an ongoing in-depth study on technical teamwork being conducted by Kiesler, Wholey and Carley (see Kiesler *et al.* (1994), Wholey, Kiesler and Carley (1995), and Carley, Kiesler and Wholey (1993), for additional coding and data collection

¹¹ When 'x' is set equal to 50 per cent of the team members for which you have maps then it is possible to place a probability distribution about the team map and calculate which team member's maps are significantly different from the team map (Banks and Carley, 1994).

details). Various types of data are collected using survey instruments, observation, and archival means. The texts I use, the answers to the two open-ended questions, are only a small set of the data collected on these students which also includes grades, grade point average, effort, and answers to several multiple page questionnaires on various communication and network issues. The instructor did not look at the answers to any of the questionnaires before assigning grades. The texts have not been analyzed by the instructor.

These two open-ended questions 'What is an information system?' and 'What leads to information system success and failure?' are interesting to look at from a mental model perspective. The first question asks the students to report on largely declarative knowledge that can be learned predominantly by memorizing definitions in various publications. The second question asks the students to report on both declarative and procedural knowledge that is learned both through experience and hearsay. From a methodological perspective it is important to know whether the effect of the coding choices depends on question type.

Coding

The texts were coded as cognitive maps using the automap procedure previously described. There were two texts for each student, one for each question. Prior to extracting the cognitive maps from these texts the following coding decisions were made: (1) Pronouns were converted to the noun that they referred to, if possible. (2) The words people or person were converted to the type of individual they referred to—user or group member. (3) For the question, 'What is an information system?' if the term information system was only implied it was added to the text. For the question 'What leads to information system success and failure?' if the terms success and failure were only implied then they were added to the text.

In order to demonstrate the impact of the critical coding choices—filtering and windowing—individual cognitive maps were extracted from these texts in several different ways. Maps were extracted from both the unfiltered and filtered texts. Additionally, maps were extracted using five different window sizes—4, 6, 8, 10 and 100. Thus, from each text 10 different cognitive maps were extracted. For each coding rule (e.g. filtered and window size of 4) a team map is constructed. The team maps are created in three ways: as the union, as a lossy intersection at degree 2 (those statements shared by at least two team members), and as a lossy intersection at degree 3 (those statements shared by at least three team members). Since teams have between four and six members these values represent a requisite agreement by 33 per cent to 75 per cent of the team.

In filtering the texts, concepts were both deleted and generalized. For each of the two questions a separate delete list and thesaurus was constructed. The delete list contained articles, conjunctions and words not contributing to the answer to that specific question. The thesaurus contained two types of generalizations: (A) Concepts which in this context had similar meanings were recoded into the same concept. For example, information system and system were both coded as information system. (B) Concepts with the same base and different tenses or endings were recoded into the same concept. For example, retrieves, retrieve, retrieving, and retrieval were all coded as retrieve.

These files were used by automap to create a filtered text from each of the unfiltered texts. Across the texts, a total of 865 unique concepts were used by the respondents; i.e. the number of concepts in the union file was 865. After filtering, the number of unique concepts dropped to 167. The number of concepts is independent of the window size.

Illustrative example

In this analysis, the following two questions are addressed. First, what is the impact of the two key coding choices—filtering and window size? And, second, is there a systematic difference in the team mental models of the successful and unsuccessful groups that is observable by examining the team maps? The related methodological question is, if there is an observable difference, is it robust across coding choices? The results of these analyses speak more to the methodological than the theoretical issues. They illustrate the potential impact of coding choices. Further, these analyses suggest the types of analyses that can be done using the procedures described for coding individual cognitive maps and deriving team maps. The results are from only eight student project groups at one point in time and so care should be taken in drawing any generalizations about the relation between team mental models and team performance from these findings. These results are simply illustrative of the types of results possible when using these methods.

Impact of coding choices

The number of concepts per map depends on the level of filtering choices. Filtered texts always have fewer concepts than do unfiltered texts (see Table 1). This is due both to deletion and generalization of concepts. Importantly, filtering the texts can also decrease the variance in the number of concepts per map. For example, in this data, there is quite a range in the size of the individuals' maps. The size of the maps does not differ significantly by question, either before or after filtering. Filtering, however, significantly decreases the size of the map (*t*-tests on the differences in means are significant at greater than the 0.0005 level for both questions). When the texts are not filtered, the number of concepts range from 7 to 90 per map. After filtering, the number of concepts per map range from 4 to 24. In general, larger maps are reduced proportionally more because they tend to contain more repetitions, more articles, and more fillers such as 'well, then', and 'let's see'. In this way, deleting words is one way of discounting for 'verbosity'. Now, let us refer to those texts that start out with more words as those being produced by more verbal students. Being verbal, at least in these texts, appears to be related to producing more concepts. That is, even after filtering, the same texts appeared to be larger and smaller. The

Table 1. Characteristics of individuals' cognitive maps with window size equal to 100

	What is an information system?		What leads to information system success and failure?	
	Average number of concepts per map	Average number of statements per map	Average number of concepts per map	Average number of statements per map
Unfiltered	32.73 (16.52)	1312.51 (1400.89)	31.07 (14.41)	1149.00 (1053.63)
Filtered	12.22 (4.25)	165.68 (105.78)	12.37 (5.07)	174.41 (141.27)
<i>t</i> -test	<i>t</i> = 7.70 <i>p</i> > 0.0005	<i>t</i> = 5.23 <i>p</i> > 0.0005	<i>t</i> = 7.84 <i>p</i> > 0.0005	<i>t</i> = 5.87 <i>p</i> > 0.0005

Standard deviations are in parentheses. *N* = 41.

Table 2. Impact of window size team cognitive maps based on filtered texts

Window Size	4	6	8	10	100
What is an information system?					
Number of unique statements (union)	305.00	401.75	469.00	513.25	625.25
Number of statements shared by two or more (lossy intersection—degree 2)	49.75	71.50	83.25	94.00	115.25
Number of statements shared by three or more (lossy intersection—degree 3)	14.25	20.75	26.00	28.75	35.50
What leads to information system success and failure?					
Number of unique statements (union)	323.12	431.13	509.88	569.92	734.62
Number of statements shared by two or more (lossy intersection—degree 2)	33.87	48.38	61.63	70.13	86.12
Number of statements shared by three or more (lossy intersection—degree 3)	8.75	11.50	14.25	16.75	22.00

correlation between the number of words in the original texts and the number of unique concepts in the filtered texts is 0.99.

The number of statements per map, depends on the size of the coding window and the filtering. There will be significantly fewer statements per map when texts are filtered (see Table 1) (*t*-tests on the differences in means are significant at greater than the 0.0005 level for both questions). There will be more statements per map the larger the size of the coding window (see Table 2). Filtering the texts can also decrease the variance in the number of statements per map. For example, in this data, there is quite a range in the complexity of the individuals' maps. When the texts are not filtered, the number of statements range from 43 to 7625 per map; after filtering, this reduces to 13 to 575 statements per map.

The number of statements and the amount of sharing increases with the size of the coding window. This is true whether or not the texts are filtered and regardless of the question. In Table 2, the impact of window size on the coding of the filtered texts is shown by degree of sharing. Interestingly, the number of shared statements (statements in the lossy intersection) typically increases slightly faster than the total number of statements (union) as window size increases. This suggests that the larger the window size the more likely it will be for the researcher to locate a degree of sharing that is higher than that which would be expected by chance. Although the number of statements in the team map increases with window size, the increase is not linear (regardless of whether a union or an intersection approach is taken). Indeed, the rate of increase decreases as the size of the window increases. For texts about a single topic, this decrease may be caused by increasing replication of common themes in the texts. That is, if the topic is small, there is only so much that can be said, and increasing the window size, after a while starts picking up this replication of ideas rather than new ideas. Whether this same diminishing returns with the size of the coding window will be seen in other contexts is a point for further research. These same trends, increasing window size increases the number of statements but with diminishing returns, hold for both types of questions. However, the diminishing returns affect seems to set in sooner for the more declarative (what is an information system) question. Since the answers for both types of question are roughly equal in length (Table 1) this slight difference in pattern may be due to a

higher level of semantic repetition on the part of the respondents when providing declarative knowledge.

Team mental models and performance

Do the mental models of successful and non-successful teams differ? Questions like this can be addressed using these methods. Following is an illustration of what such an analysis might look like. The results are purely illustrative and can be treated as interesting speculations to be considered in future research. Each text is analyzed both before and after filtering and for windows of size 4, 6, 8, 10, and 100. The results indicate that on average members of successful groups tend to use more concepts and have larger cognitive maps (more statements) than do members of non-successful groups both in defining an information system and in thinking through the conditions of success and failure. In Table 3, this difference for a coding window of size 100 is shown. This pattern holds for all but the very smallest window sizes, e.g. 4. The difference between successful and non-successful groups is significant when the data is unfiltered but not when the data is filtered. This suggests that although successful groups may be more verbal, they may not rely on more concepts. To be successful, groups may need to be able to express the same sentiment in multiple ways.

It has been suggested that to be effective, teams (particularly teams of managers), need to have a shared mental model (Porac and Thomas, 1990). In this data, relatively little information is shared across the individual cognitive maps. Nevertheless, there is on average greater sharing in the successful than non-successful groups (see Table 4). The team map of successful groups tends to be larger and the degree of sharing higher. Both the number of statements shared by two or more team members, and by three or more team members, is larger for successful and non-successful

Table 3. Characteristics of individuals' cognitive maps for successful and non-successful teams with window size equal to 100

	What is an information system?		What leads to information system success and failure?	
	Average number of concepts per map	Average number of statements per map	Average number of concepts per map	Average number of statements per map
Successful				
Unfiltered	36.44 (18.82)	1633.88 (1664.81)	34.52 (14.77)	1378.28 (1162.60)
Filtered	12.64 (4.49)	178.52 (112.70)	13.28 (4.92)	196.20 (153.33)
Non-successful				
Unfiltered	26.94 (10.13)	810.38 (598.08)	25.69 (12.39)	790.75 (757.10)
Filtered	11.56 (3.90)	145.63 (93.87)	10.94 (5.13)	140.38 (116.54)
<i>t</i> -test				
Unfiltered	<i>t</i> = 2.09 <i>p</i> > 0.05	<i>t</i> = 2.26 <i>p</i> > 0.025	<i>t</i> = 2.06 <i>p</i> > 0.05	<i>t</i> = 1.96 <i>p</i> > 0.05
Filtered	<i>t</i> = 0.81 <i>p</i> > 0.25	<i>t</i> = 1.10 <i>p</i> > 0.25	<i>t</i> = 1.45 <i>p</i> > 0.10	<i>t</i> = 1.32 <i>p</i> > 0.25

Standard deviations are in parentheses. *N* = 25 for successful and 16 for non-successful.

groups. (The one exception here is when an extremely small window, size 4, is used.) Finally, members of successful groups bring more ideas to the table that can draw on than do non-successful groups. We see this in that the number of unique statements (union map) available to the group is larger for successful than non-successful groups. The differences in Table 4, due to the small number of groups, are not significant. However, the pattern is consistent with the claim of cognitive similarity breeding success.

The team mental model however, is expected to be more than the collection of dyadic models. As previously noted, we can derive a team map by extracting those statements that occur in the maps of the majority of the team members' maps. For most teams examined, that would be all

Table 4. Impact of window size team cognitive maps for successful and non-successful teams using filtered texts

Window Size	4	6	8	10	100
Successful					
What is an information system?					
Number of unique statements (union)	311.60	408.00	476.00	523.60	652.00
Number of statements shared by two or more (lossy intersection—degree 2)	53.29	74.40	88.00	98.00	123.20
Number of statements shared by three or more (lossy intersection—degree 3)	13.60	21.20	27.60	31.60	40.00
What leads to information system success and failure?					
Number of unique statements (union)	345.80	460.20	549.00	614.20	808.20
Number of statements shared by two or more (lossy intersection—degree 2)	34.20	50.60	63.40	73.40	91.40
Number of statements shared by three or more (lossy intersection—degree 3)	11.20	14.40	16.80	20.00	24.40
Non-successful					
What is an information system?					
Number of unique statements (union)	296.00	391.33	456.67	496.00	580.67
Number of statements shared by two or more (lossy intersection—degree 2)	44.00	66.67	75.33	86.00	102.00
Number of statements shared by three or more (lossy intersection—degree 3)	15.33	20.00	23.33	24.00	28.00
What leads to information system success and failure?					
Number of unique statements (union)	285.33	382.67	444.67	495.33	612.00
Number of statements shared by two or more (lossy intersection—degree 2)	33.33	44.67	58.67	64.67	77.33
Number of statements shared by three or more (lossy intersection—degree 3)	4.67	6.67	10.00	11.33	18.00

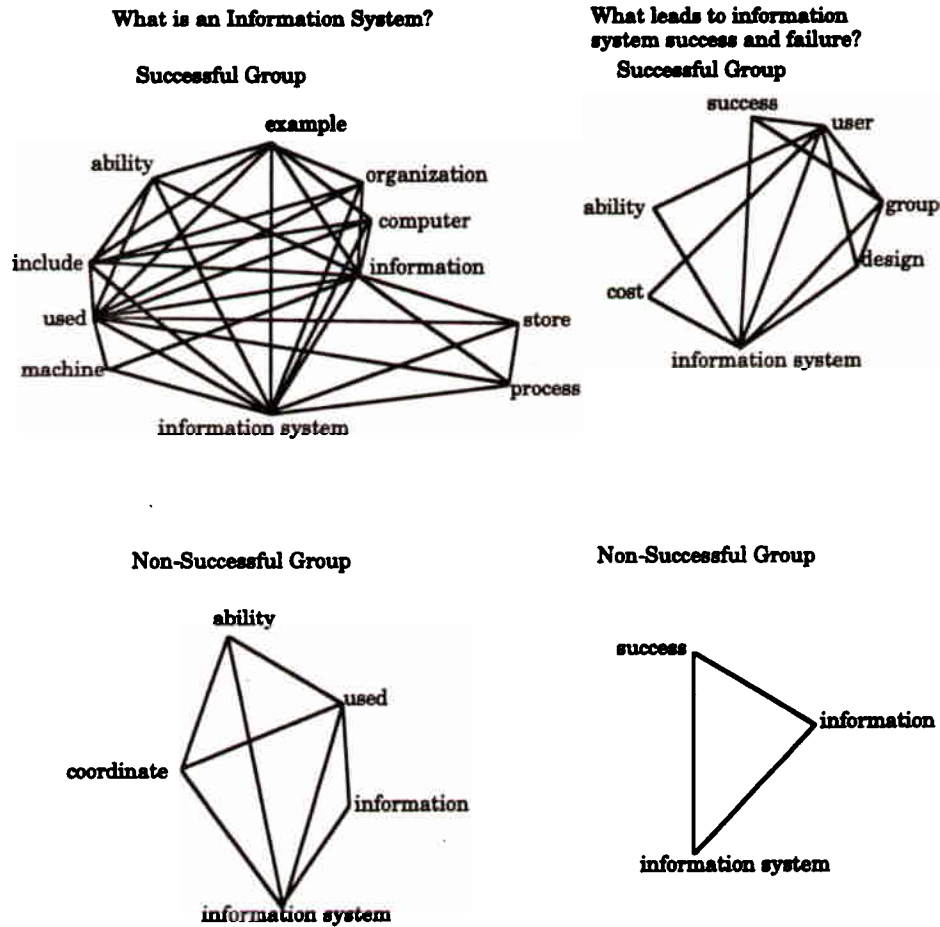


Figure 3. Illustrative team maps

statements in three or more of the team members' maps. Illustrative team maps are shown in Figure 3 for a successful and non-successful group. Note, that the successful group's team map is both more elaborate than, and contains different information from, the non-successful group. For example, the successful team sees project success as being a function of finding a design for the information system that meets the user's abilities and cost requirements. Whereas, collectively, the non-successful group sees success as hinged on the ability of the information system to provide information.

Managers and team mental models

Much of the research on managerial mental models has focused on the extent to which manager's models are homogeneous with each other (Reger, 1990; Reger and Huff, 1993; Daniels *et al.*, 1994). However, we can also ask, to what extent are manager's cognitions similar to team members? This type of question can also be addressed using these methods. Following is an illustration of what such an analysis might look like. Again, the results are purely illustrative and can be treated as interesting speculations to be considered in future research.

There are many reasons to expect that managers will have models similar to those of their team members. Specifically, interaction tends to lead to similar mental models (Carley, 1991), manager's mental models are affected by their local social and cultural environment (Huff, 1982), and each organization may have its own culture which is reflected in the team and manager's mental models (Johnson, 1987). Klein and Posey (1986) suggested that effective supervision required the manager to share in the team's mental model. Alternatively, Larson and Christensen (1993) suggest that effectiveness may be due to the manager recognizing how other's mental models compliment each other. In this case, little overlap is expected between the manager's map and that of the team members. One way of addressing these hypotheses is by seeing whether, at the dyad level, managers of successful groups share more information with team members than do managers of non-successful groups. Previous research has suggested that experts hold more elaborate and detailed mental models (Murphy and Wright, 1984; Tanaka and Taylor, 1991). Thus, before addressing the question of managerial supervision, it is necessary to determine whether or not managers (especially those of successful groups), as experts, have more elaborate maps. If they do, then they may share more with team members simply because of the greater complexity of their maps. Managers, however, do not hold more elaborate maps than other team members on average (they are not the most verbal students and the number of concepts/statements they have in their maps is not significantly different from the number of concepts/statements held by other team members).

Now let us turn to the difference between successful and non-successful groups. For this data, the hypothesis that in successful groups the manager's map will be more similar to the team members, does not hold for the first question (What is an information system?) but it does hold for the second question (What leads to information system success and failure?). For question 1, the managers of the successful groups share an average of 4.67 concepts and 20.39 statements with each of the other group members; whereas, the managers of the non-successful groups share an average of 4.78 concepts and 20.63 statements with each of the other team members. In contrast, for question 2, the managers of successful groups share an average of 4.59 concepts and 21.66 statements with each of the other group members; whereas, the managers of non-successful groups share an average of 3.67 concepts and 13.93 statements with each of the other team members¹². This difference is even more surprising given that, on average, the successful managers' maps are near in size to those of the other group members than are the non-successful managers' maps. An avenue for future exploration is whether and how this relative difference contributes to performance.

These results suggest that what information the manager shares with the group members may be more important than how much information the manager shares with the team. In this case, a common view of factors affecting performance (procedural knowledge) appears more important than a common definition of the subject matter. In fact recognizing differences in declarative knowledge may be important for success. This conclusion must be viewed with caution given the small number of teams and the fact that managerial recognition of differences was not measured directly. Nevertheless, these results do suggest that this methodological approach for examining issues of managerial cognition by contrasting individual and team maps is promising.

¹² Another way of addressing this issue would be to compare the Hamming distance (Hamming, 1950, 1980) between the manager's map and the team map and the Hamming distance of other's cognitive maps to the team map (Banks and Carley, 1994). Since each map is a binary graph, the Hamming distance between two maps is a measure of overall similarity. In general, statistical techniques for analyzing networks in their entirety are quite primitive; however, such techniques are needed for analyzing cognitive maps and the similarity between them. Most researchers analyzing cognitive maps do not statistically address the similarity of maps directly, see for example Daniels *et al.* (1994).

Discussion

A technique, called map analysis, for coding and comparing abstractions of individual mental models as cognitive maps and generating team maps was presented. This technique has the advantage that the researcher can decide what types of information or knowledge are appropriate to consider as possible candidates for information in the maps retrospectively. In other words, this technique does not place *a priori* requirements on map content. Rather, the map can be coded automatically from any textual source. Another advantage to map analysis is that this technique can be used recursively. That is, imagine a hierarchy of teams, such that team AB is composed of team A and team B and each of these teams is in turn composed of multiple individuals. Map analysis allows the researcher to first generate the team map for teams A and B separately, and then to find the combined team map for the larger team AB. The same techniques that are useful for comparing and contrasting individuals' cognitive maps can be used to compare and contrast teams' maps. Another advantage is that this technique allows the researcher to extract the team map, the differences between each team member's map and the team map, and the extent to which the maps of pairs or groups of individuals may be similar to or distinct from the team map.

The foregoing discussion assumed that the researcher had access to a series of texts from which individual cognitive maps and consequently teams maps could be extracted. In this paper the impact of the data collection technique on the results was not examined. There are many ways in which data collection might affect the results obtained from a textually based cognitive mapping technique. Let us consider two of these. First, the collection technique affects the level of detail in the coded map. In general, better results are achieved from texts that focus on a belief, action, process, and/or definition that is relevant to the task(s) the team is engaged in. Techniques that encourage the team members to provide more detail should be used; however, caution should be taken in the analysis as issues of researcher reactivity arise. Such techniques include: prompting, multiple questions on the same topic (triangulation), and repeated questions (in an interview setting). Additional research is needed on the most effective technique for eliciting these texts. Second, collection techniques affect the interpretation of the results. As noted there is some disagreement over whether or not team members need to be aware that they share information for it to be in the team model. Different collection techniques are needed to collect team data where members are aware that they share information and where they are not necessarily aware that they share information. In this paper, the data was collected in such a way that it was not known whether or not the team members knew what information they shared. Presumably, a data collection technique that ensured that the texts contained information on what the team members knew to be shared would result in more shared concepts and statements, with little effect due to degree of consensus. Whether or not there would be interactions with other coding choices is a point for further research.

In the illustrative analysis the potential impact of several critical coding decisions were examined. These decisions centered around filtering, windowing and the requisite degree of sharing for something to be considered part of the team map. Filtering involves making decisions about which concepts to place in the deletion list and which concepts to generalize in the thesaurus list. Filtering leads to smaller more parsimonious maps and focuses attention on the common issues. Examining differences for both filtered and unfiltered text can provide insight into the relative impact of verbal variety as opposed to conceptual variety. However, filtering has the potential to introduce research bias. Windowing involves choosing the maximal distance between concepts in which a relationship can exist. Increasing window size can increase the number of statements in a map, i.e. the degree of interconnectivity among concepts. Consequently,

increasing window size may increase the likelihood that similarities in the way in which concepts are linked will emerge. On the down side, to small a window size can miss author intended links (under-connecting) and extremely large windows are likely to create links that the author never intended (over-connecting). Finally, increasing the degree of consensus required for a statement to be in the team map decreases the team map's complexity; but, it increases the likelihood that team members will recognize that the statement is shared. Researchers need to be aware how the coding choices they make will affect the results (see also Carley, 1993). In conducting this type of analysis, researchers must weigh the tradeoffs involved in specific types of filtering, and window size, a sharing criterion. Importantly, the map analysis tools are flexible enough to be used with any choice about filtering, window size, or degree of consensus.

Neither filtering, altering the window size, nor altering the degree of consensus guarantees that significant differences will emerge. Many differences will be robust across these coding decisions. For example, the qualitative difference between team maps (such as the difference between successful and non-successful teams) is not affected by most changes in window size, except for very small windows; however, whether or not the difference is statistically significant may be dependent on window size. Additionally, the results appear to be more affected by filtering and by the definition of what is included (union or lossy intersection) than by window size. Further research is needed on whether or not it makes sense to develop standard procedures for word deletion, generalization, a standard window size choice, and a standard degree of consensus for creating team maps.

The proposed approach to team mental models does not rely on the individuals in the team exhibiting conscious awareness that the team mental model is shared by team members. Whether or not team members have a perception that the derived team map is shared by the team is something that the researcher may or may not wish to determine. These techniques can be used to derive maps that represent team mental models, and to test the extent to which an espoused team model that is perceived as being shared is actually shared by the members of the team. Additionally, this approach is agnostic as to how much consensus is required for something to be considered part of the team model or shared culture. These techniques can be used to examine different degrees of sharing and to see how this alters the results.

In this paper, an automated approach to coding that placed links between concepts that were within so many positions of each other was used. This procedure is embodied in the program *automap*. This automated approach is more cost-effective than a semi-automated one where the researcher decides which pairs of concepts form a statement, as the automated approach reduces coder time and obviates the need to do inter-rater reliability checks. One issue is how much time is saved by using these automated techniques. Typically, reading and coding a text of 1 to 4 pages may take anywhere from 1 to 5 hours to code depending on the complexity of the coding scheme. With this approach, coding a text takes only a few minutes. Further, time is saved in not having to train human coders or run inter-rater reliability checks. Another 'cost saving' advantage of this approach has to do with ease of re-analysis. For example, imagine that the researcher, the first time the data was analyzed, had failed to extract a particular concept that was later deemed relevant. Or, imagine that the researcher realizes after having coding most of the texts that there is not really a difference in meaning in the way two different concepts are used. Using classic content analysis techniques, both the need for greater detail or the need to combine old categories might necessitate going back and re-reading all texts and recoding those portions associated with the change. In contrast, for the approach just demonstrated, the re-analysis would require simply changing the thesaurus or delete list and then re-running the program for extracting maps. This would be a matter of a few minutes not hours. These cost savings make possible the coding of greater quantities of texts and admit more thorough empirical analysis. This is not to suggest that

the automated approach is without costs. There is still a substantial time effort in collecting the texts, transcribing them, cleaning them to eliminate spelling errors, creating the concept lists, and generating the computer script files for running the analyses. Nevertheless, the larger the number of texts and the more complex the coding scheme the greater the possible savings that can be realized using an automated approach.

A second issue is whether the resultant maps from an automated approach are as accurate as those generated by the researcher using a semi-automated approach. This is a very important point on which additional research is needed. Techniques for handling metaphors, homonyms, and colloquialisms are critical in textual analysis. An automated approach may inadvertently encourage the researcher to ignore differences in meaning and to simply focus on default meanings attributed to words. This may lead to an over-representation of certain relations and an under-representation of others.

Automated approaches are not without limitations. For example, Namenwirth and Weber (1987) argue that automated content analysis is most useful for analyzing large trends and not individual cognition. However, cognitive mapping techniques move beyond content analysis and make possible the analysis of mental models. Huff (1990) has argued that automated and simple coding is best for scanning large amounts of data and looking for anomalies which can then be explored in depth using more detailed, labor-intensive techniques. However, advances in automated textual analysis are reducing the validity of this claim by increasing the level of automation possible even for more detailed coding approaches (Roberts, 1997). Despite the enormous progress that has been made in the past decade, automated tools for coding texts are still relatively limited. For example, the tool used here does not pick up various nuances such as the use of metaphor, the direction of the relation, or semantic cues about the strength of the relation. To get at such factors more costly labor-intensive techniques that might require more judgments are needed.

Using map analysis, the mental models of the members of eight different project teams and a total of 41 individuals as they began the process of project development were explored. This study of team behavior serves simply to illustrate the value of an automated approach to looking at team cognition and the impact of coding choices on the resultant analysis. Nevertheless, the results are interesting and suggest possible hypotheses for future studies. The results suggest that the mental models of successful and non-successful teams do differ systematically at the start of the project. In particular, for successful groups, members tend to have more elaborated maps, members appear to be more verbal, and team maps are more detailed than for non-successful groups. Since competency was controlled for when the teams were formed, these results cannot simply be attributed to successful groups being initially intellectually or experientially advantaged. Whether these results hold for other types of teams, whether successful teams maintain common mental models over time, and whether the team mental model that emerges differs systematically from that of non-successful teams are all interesting avenues for further research.

An advantage of these techniques for analyzing texts as cognitive maps is that they allow the researcher to take rich verbal data and analyze it empirically. This makes it possible to statistically test hypotheses about the formation, maintenance, and change in team mental models over time and across teams. These techniques make it possible to address new and interesting questions in the areas of performance, motivation, and organizational learning. For example, learning can be examined by looking at additions to and deletions from older maps, in addition to looking at specific content. Organizational learning can be examined by looking at the creation and change in the team map. Issues of leadership can be addressed by contrasting the manager's map with that of other team members or the team map. Indeed, the procedures for extracting, contrasting, and combining maps are quite general and have a number of applications in the

organizational field. Interesting applications include the following: comparisons of expert and novice mental models, explorations of team mental models, studies of organizational culture in terms of differences in individuals' and teams' mental models, and examinations of organizational fields. Importantly, these techniques allow the researcher to capture information on the degree of change or difference (number of concepts/statements) and the content (what the concepts/statements are). By examining both degree and content a more complete understanding of issues involving team mental models may emerge. The techniques presented herein can provide the basis for a consistent and theoretically grounded approach to the quantitative study of team mental models.

References

- Anderson, J. R. and Bower, G. H. (1973). *Human Associative Memory*, Winston-Wiley, Washington DC.
- Axelrod, R. M. (1976). *Structure of Decision: The Cognitive Maps of Political Elites*, Princeton University Press, Princeton, NJ.
- Banks, D. and Carley, K. M. (1994). 'Metric inference for social networks', *Journal of Classification*, **11**, 121–149.
- Bannister, D. and Mair, J. M. M. (1968). *The Evaluation of Personal Constructs*, Academic Press, London.
- Barley, S. (1983). 'Semiotics and the study of occupational and organizational cultures', *Administrative Science Quarterly*, **28**, 393–413.
- Blumer, H. (1969). *Symbolic Interactionism: Perspective and Method*, Prentice Hall, Englewood Cliffs, NJ.
- Bobrow, D. G. and Norman, D. A. (1976). 'Some principles of memory schemata'. In: Bobrow D. G. and Collins, A. (Eds) *Representation and Understanding: Studies in Cognitive Science*, Academic Press, New York, NY.
- Carley, K. M. (1984). 'Constructing consensus'. Unpublished doctoral dissertation, Harvard.
- Carley, K. (1986a). 'An approach for relating social structure to cognitive structure', *Journal of Mathematical Sociology*, **12**(2), 137–189.
- Carley, K. (1986b). 'Knowledge acquisition as a social phenomenon', *Instructional Science*, **14**, 381–438.
- Carley, K. (1988). 'Formalizing the social expert's knowledge', *Sociological Methods and Research*, **17**, 165–232.
- Carley, K. M. (1991). 'A theory of group stability', *American Sociological Review*, **56**(3), 331–354.
- Carley, K. M. (1993). 'Coding choices for textual analysis: A comparison of content analysis and map analysis'. In: Marsden, P. (Ed.) *Sociological Methodology*, Vol. 23, Blackwell, Oxford, pp. 75–126.
- Carley, K. M. (1997). 'Network text analysis: The network position of concepts'. In: Roberts, C. W. (Ed.) *Text Analysis for the Social Sciences: Methods for Drawing Statistical Inferences from Texts and Transcripts*, Lawrence Erlbaum Associates, Mahwah, NJ, pp. 79–100.
- Carley, K. M. and Kaufer, D. S. (1993). 'Semantic connectivity: An approach for analyzing semantic networks', *Communication Theory*, **3**(3), 183–213.
- Carley, K. M. and Palmquist, M. (1992). 'Extracting, representing and analyzing mental models', *Social Forces*, **70**, 601–636.
- Carley, K., Kiesler, S. and Wholey, D. (1993). 'Learning teamwork: studies of training in software development'. In: *Proceedings of the 1993 Coordination Theory and Collaboration Technology Workshop*. Symposium conducted for the National Science Foundation, Washington, D.C.
- Charniak, E. (1972). 'Toward a model of children's story comprehension'. Doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA.
- Clark, H. H. and Clark, E. V. (1977). *Psychology and Language: An Introduction to Psycholinguistics*, Harcourt Brace Jovanovich, New York, NY.
- Collins, A. M. and Loftus, E. P. (1975). 'A spreading-activation theory of semantic processing', *Psychological Review*, **82**, 407–428.
- Daniels, K., Johnson, G. and de Chernatony, L. (1994). 'Differences in managerial cognitions of competition', *British Journal of Management*, **5**, 21–29.

- Danowski, J. A. (1982). 'A network-based content analysis methodology for computer-mediated communication: an illustration with a computer bulletin board'. In: *Communication Yearbook*, Sage, New York, pp. 904–925.
- Danowski, J. A. (1988). 'Organizational infographics and automated auditing: using computers to unobtrusively gather and analyze communication'. In: Goldhaber, G. and Barnett, G. (Eds) *Handbook of Organizational Communication*, Ablex, Norwood, NJ, pp. 385–433.
- Durkheim, E. (1938). *Rules of Sociological Method*. Free Press, New York.
- Eden, C., Jones, S. and Sims, D. (1979). *Thinking in Organizations*, Macmillan Press, London, England.
- Edwards, D. and Mercer, N. (1986). 'Context and continuity: classroom discourse and the development of shared knowledge'. In: Durkin, K. (Ed.) *Language Development in the School Years*, Brookline Books, Cambridge, MA, pp. 172–202.
- Ericsson, K. A. and Simon, H. A. (1984). *Protocol Analysis: Verbal Reports as Data*, MIT Press, Cambridge, MA.
- Fauconnier, G. (1985). *Mental Spaces: Aspects of Meaning Construction in Natural Language*, Bradford Books, MIT Press, Cambridge, MA.
- Fiol, C. M. (1994). 'Consensus, diversity, and learning in organizations', *Organization Science*, 5(3), 403–420.
- Forman, E. A., Minick, N. and Stone, C. A. (Eds) (1993). *Contexts for Learning: Sociocultural Dynamics in Children's Development*, Oxford University Press, New York.
- Franzosi, R. (1990a). 'From words to numbers: a generalized and linguistic-based coding procedure for collecting textual data'. In: Clogg, C. C. (Ed.) *Sociological Methodology*, Jossey-Bass, San Francisco, CA.
- Franzosi, R. (1990b). 'Computer-assisted coding of textual data: an application to semantic grammars', *Sociological Methods and Research*, 19, 225–257.
- Goffman, E. (1974). *Frame Analysis: An Essay on the Organization of Experience*, Harper and Row, New York, NY.
- Hamming, R. W. (1950). 'Error detecting and error correcting codes', *Bell System Technical Journal*, 29, 147–160.
- Hamming, R. W. (1980). *Coding and Information Theory*, Prentice-Hall, Englewood Cliffs, NJ.
- Hergenhahn, B. R. (1994). 'Psychology's cognitive revolution', *American Psychologist*, 49(9), 816–817.
- Huff, A. S. (1982). 'Industry influences on strategy formulation', *Strategic Management Journal*, 3, 119–130.
- Huff, A. S. (1990). *Mapping Strategic Thought*, Wiley, Chichester.
- Hutchins, E. (1990). 'The technique of team navigation'. In: Galegher, J., Kraut, R. and Egido, C. (Eds) *Intellectual Teamwork*, Lawrence Erlbaum Associates, Hillsdale, NJ.
- Hutchins, E. (1991a). 'Organizing work by adaptation', *Organizational Science*, 2, 14–39.
- Hutchins, E. (1991b). 'The social organization of distributed cognition'. In: Resnick, L. B., Levine, J. M. and Teasley, S. D. (Eds) *Perspectives on Socially Shared Cognition*, American Psychological Association, Washington D.C.
- Johnson, G. (1987). *Strategic Change and the Management Process*, Blackwell, Oxford.
- Johnson-Laird, P. N. (1983). *Mental Models: Toward a Cognitive Science of Language, Inference, and Consciousness*, Harvard University Press, Cambridge, MA.
- Kaufer, D. S. and Carley, K. M. (1993). *Communication at a Distance: The Effect of Print on Socio-Cultural Organization and Change*, Lawrence Erlbaum Associates, Hillsdale, NJ.
- Kelly, G. A. (1955). *The Psychology of Personal Constructs*, Vols 1 and 2, Norton, New York, NY.
- Kiesler, S., Wholey, D. and Carley, K. (1994). 'Coordination as "linkage": the case of software development teams'. In: Harris, D. H. (Ed.) *Organizational Linkages: Understanding the Productivity Paradox*, National Academy Press, Washington, D. C., pp. 214–239.
- Klein, J. A. and Posey, P. A. (1986). 'Good supervisors are good supervisors—anywhere', *Harvard Business Review*, 64(6), 125–128.
- Klimoski, R. and Mohammed, S. (1994). 'Team mental model: construct or metaphor?' *Journal of Management*, 20(2), 403–437.
- Knorr-Cetina, K. (1981). *The Manufacture of Knowledge*, Pergamon, Oxford, England.
- Krackhardt, D. (1987). 'Cognitive social structures', *Social Networks*, 9, 109–134.
- Larson, J. R. and Christensen, C. (1993). 'Groups as problem solving units: towards a new meaning of social cognition', *British Journal of Social Psychology*, 32, 5–30.
- Latour, B. and Woolgar, S. (1979). *Laboratory Life*, Sage, Beverly Hills, CA.

- Martin, J. (1992). *Cultures in Organizations: Three Perspectives*, Oxford University Press, New York, NY.
- Mead, G. H. ([1934]1962). *Mind, Self, and Society*, University of Chicago Press, Chicago, IL.
- Minsky, M. A. (1975). 'A framework for representing knowledge'. In: Winston, P. (Ed.) *The Psychology of Computer Vision*, McGraw-Hill, New York, NY.
- Murphy, G. L. and Wright, J. C. (1984). 'Changes in the conceptual structure with expertise: differences between real-world experts and novices', *Journal of Experimental Psychology: Learning, Memory and Cognition*, **10**, 144–155.
- Namenwirth, J. Z. and Weber, R. P. (1987). *Dynamics of Culture*, Allen & Unwin, Boston, MA.
- Newell, A. and Simon, H. A. (1972). *Human Problem Solving*, Prentice-Hall, Englewood Cliffs, NJ.
- Ogilvie, D. M., Stone, P. A. and Kelly, E. F. (1982). 'Computer-aided content analysis'. In: Smith, R. B. and Manning, P. K. (Eds) *A Handbook of Social Science Methods*, Ballinger, New York, NY.
- Polanyi, M. P. (1958[1962]). *Personal Knowledge: Towards a Post-Critical Philosophy*, The University of Chicago Press, Chicago, IL.
- Porac, J. F. and Thomas, H. (1990). 'Taxonomic mental models in competitor definition', *Academy of Management Review*, **15**, 224–240.
- Reger, R. (1990). 'Management thought structures and competitive positioning'. In: Huff, A. (Ed.) *Management Thought Structures and Competitive Positioning*, Wiley, New York, NY.
- Reger, R. K. and Huff, A. S. (1993). 'Strategic groups: a cognitive perspective', *Strategic Management Journal*, **14**, 103–124.
- Reitman, J. S. and Rueter, H. R. (1980). 'Organization revealed by recall orders and confirmed by pauses', *Cognitive Psychology*, **12**, 554–581.
- Roberts, C. W. (1989). 'Other than counting words: a linguistic approach to content analysis', *Social Forces*, **68**, 147–77.
- Roberts, C. W. (Ed.) (1997). *Text Analysis for the Social Sciences: Methods for Drawing Statistical Inferences from Texts and Transcripts*, Lawrence Erlbaum Associates, Mahwah, NJ.
- Romney, A. K., Weller, S. C. and Batchelder, W. H. (1986). 'Culture as consensus: A theory of culture and informant accuracy', *American Anthropologist*, **88**(2), 313–338.
- Rumelhart, D. E. and Ortney, A. (1976). 'The representation of knowledge in memory'. In: Anderson, R. C., Spiro, R., Spiro, J. and Montague, W. E. (Eds) *Schooling and the Acquisition of Knowledge*, Lawrence Erlbaum Associates, Hillsdale, NJ.
- Schank, R. C. and Abelson, R. P. (1977). *Scripts Plans and Goals and Understanding*, Wiley, New York, NY.
- Schank, R. C. and Colby, K. M. (1973). *Computer Models of Thought and Language*, W. H. Freeman, San Francisco, CA.
- Shavelson, R. J. (1972). 'Some aspects of the correspondence between content structure and cognitive structure in physics instruction', *Journal of Educational Psychology*, **63**, 225–234.
- Sowa, J. F. (1984). *Conceptual Structures*, Addison-Wesley, Reading, MA.
- Sperry, R. W. (1995). 'The riddle of consciousness and the changing scientific worldview', *Journal of Humanistic Psychology*, **35**(2), 7–33.
- Stone, P. J. and Cambridge Computer Associates (1968a). *User's Manual for the General Inquirer*, MIT Press, Cambridge, MA.
- Stone, P. J., Dunphy, D. C., Smith, M. S. and Ogilvie, D. M. (1968b). *The General Inquirer: A Computer Approach to Content Analysis*, MIT Press, Cambridge, MA.
- Stryker, S. (1980). *Symbolic Interactionism*, Benjamin Cummings, Menlo Park, CA.
- Tanaka, J. W. and Taylor, M. (1991). 'Object categories and expertise: Is the basic level in the eye of the beholder?' *Cognitive Psychology*, **23**, 457–482.
- VanLehn, K. and Brown, J. S. (1980). 'Planning nets: a representation of formalizing analogies and semantic models of procedural skills'. In: Snow, R. E., Frederico, P. A. and Montague, W. E. (Eds) *Aptitude Learning and Instruction: Vol. 2, Cognitive Process Analysis and Problem Solving*, Lawrence Erlbaum Associates, Hillsdale, NJ.
- VanLehn, K. and Garlick, S. (1987). 'Cirrus: An automated protocol analysis tool'. In: Langley, P. (Ed.) *Proceedings of the Fourth International Workshop on Machine Learning*, Morgan-Kaufman, Los Altos, CA.
- Walsh, J. P. (1995). 'Managerial and organizational cognition: notes from a trip down memory lane', *Organization Science*, **6**(3), 280–321.
- Weick, K. E. and Roberts, K. A. (1993). 'Collective mind in organizations: heedful interrelating on flight decks', *Administrative Science Quarterly*, **38**, 357–381.

- Weitzman, E. A. and Miles, M. B. (1995). *Computer Programs for Qualitative Data Analysis*, Sage, London, England.
- Wholey, D., Kiesler, S. and Carley, K. M. (1995). 'Learning teamwork: a study of communication, structure, and success in novice software development teams'. Working paper, Social and Decision Sciences Department, Carnegie Mellon University.
- Winston, P. H. (1977). *Artificial Intelligence*, Addison-Wesley, Reading, MA.
- Wyer, R. S. Jr. and Carlston, D. E. (1979). *Social Cognition, Inference, and Attribution*, Lawrence Erlbaum Associates, Hillsdale, NJ.