Order Number 9108088

Mental representations of organizational structure: An investigation of the psychological underpinnings of the division of labor

Clinchy, Ross McVicker, Ph.D.

City University of New York, 1990

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106

MENTAL REPRESENTATIONS OF ORGANIZATIONAL STRUCTURE: An Investigation of the Psychological Underpinnings of the Division of Labor

by

Ross McVicker Clinchy

A dissertation submitted to
the Graduate Faculty in Psychology
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy, The City University of New York

1990

This manuscript has been read and accepted for the Graduate Faculty in Psychology in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy

5/22/90
Date

Chair of Examining Committee

May 30 1990
Date

Executive Officer

Anderson J. Franklin
Louis J. Gerstman
James S. Krantz
Kathleen Pogue White

Supervisory Committee

The City University of New York

ABSTRACT

MENTAL REPRESENTATIONS OF ORGANIZATIONAL STRUCTURE: AN INVESTIGATION OF THE PSYCHOLOGICAL UNDERPINNINGS OF THE DIVISION OF LABOR

by

Ross McVicker Clinchy

Advisor: Professor Laurence J. Gould

The two studies reported here represent an attempt to use similarity judgments to reconstruct an individual's mental representations of the structure of relationships in the social groups of which they are a member. The rationale for the method was derived from the sociotechnical theory of organizations developed at the Tavistock Institute.

Two studies were performed. In the first, ten members of the staff of an academic health center were asked to rate the similarity of all pairs of twenty five of the senior members of the center. The resulting similarity matrix was fitted to the multidimensional scaling, hierarchical clustering, additive tree, and extended tree models. The resulting output configurations bore a strong resemblance to

the formal organizational chart for the center, and the discrepancies between the reconstructions and the formal chart could be interpreted as reflecting ways in which the actual structure of relationships in the organization failed to conform to the ideal represented by the formal chart.

In the second study, similarity ratings were gathered from twelve of the members of an interdisciplinary healthcare team on all pairings of the eighteen members of the team. The resulting similarity matrix was also fitted to several of the statistical models used in the first study. The results were used to derive a set of hypotheses about the structure of relationships in the team, which were submitted to the subjects for validation and were further pursued in individual interviews. In general, the hypotheses were confirmed by the subjects, suggesting that the real structure of relationships in a group can be determined by this method.

The discussion of these studies includes a review of the procedures for applying this method and of the further studies that could be performed in order to refine the technique. It also contains a discussion of the difficulties that are encountered in attempting to draw conclusions regarding the mental representations that are presumed to govern perception and behavior in organizations.

ACKNOWLEDGMENTS

There are a number of people who I wish to thank for helping me with this dissertation. Larry Gould played a number of roles. As Chair of my Committee, he was always available to me for guidance and encouragement, and provided reassurance that there would be an end to the process and that it would be worth the effort and difficulty. At the same time, he allowed me an exceptional degree of freedom to work at my own pace and in my own way. Furthermore, he suggested a modification in the experimental design that made the study both more elegant and more feasible.

The other members of my Committee, A. J. Franklin and Louis J. Gerstman, were also consistently supportive but, at the same time, able to be critical. My work has benefitted greatly from their comments and suggestions. A. J. Franklin was especially helpful in encouraging me to improve the clarity of my exposition and to keep the many different threads in the work from becoming entangled with each other. Lou Gerstman was a valuable resource on the methodological and statistical issues, but he also encouraged me to think about the clinical implications of the work and to avoid becoming too engrossed in its technical intricacies.

Two other people made critical contributions to the development of the project reported here. It was Jim Krantz who, after I had shown him the Pilot Study data, drew the crucial and highly illuminating connection between similarity and sentience, which showed me how to put what I was doing into a theoretical context.

Jim Corter was invaluable in helping me navigate through deep statistical and methodological waters in which I probably would have foundered if I had been alone. Furthermore, it was he who made me aware of the EXTREE program that he developed, the statistical model that proved to be the most useful in analyzing the data that I collected.

I owe a special debt to my subjects. For the two subject groups who were willing to volunteer their time, there was really very little reward and, I think, no small measure of anxiety involved. I am very grateful to them for their willingness to help me.

Throughout the period in which I carried out this project, I was a member of the staff of the State University of New York's Health Science Center at Brooklyn, and I appreciate the patience and tolerance of my colleagues for the times in which I was preoccupied with the dissertation

and less available to them and to the work of the irstitution. Two of my colleagues there, Jim Ranck and Len Rosenblum, were particularly helpful in that they were willing to spend several hours going over the project and helping me understand what I had found and in thinking about how to explain it clearly and concisely.

Finally, I would like to thank my families: my parents, Evans and Blythe Clinchy; my brothers, Sean and Gavin, their wives, Estelle and Vanessa, and my nephew, Evans; and my second family, the Chanils, Leo, Lisa, Debra, Mark and Justin; and most especially, my wife, Elayne. Any work like this, which is done over the course of several years, makes enormous demands on one's family and involves considerable sacrifice. I have been very fortunate to have two families who have been unfailingly interested without being intrusive, encouraging but patient, and, above all, loving, and I am deeply grateful to them for everything that they have given me.

Table of Contents

List of Tables

List of Figures

List of Graphs

I. INTRODUCTION

In a general sense, the studies described here are an attempt to explore the psychological mechanisms that support the division of labor. Specifically, they are concerned with the mechanisms that enable an individual to take up a role in a work situation and to direct his or her behavior in accordance with the demands and constraints of that role.

Modern work organizations typically possess immense technological, material, and human resources, which must be organized and coordinated so that they are working together, in a reasonably harmonious way, to accomplish the task of the organization. As our society has become more technologically oriented, we have become increasingly inclined to see organizations as large machines, in which the structure of the organization is largely determined by the demands of the technology employed. In a factory, for example, the technical layout may define the stages of a particular production process, which, in turn, defines certain basic divisions in the organizational structure. From such a perspective, it is easy to conclude that the organization's ability to perform its task is primarily determined by the capacities and limitations of its machines.

At the same time, however, it is worth remembering that, at least so far, all work organizations still require people in order to function, and that, furthermore, they require people who are capable of doing the right thing at the right time. Without such people, even a highly technological enterprise will be inert, and it seems apparent that an organization's ability to perform its task effectively will depend, in part, on the degree to which its human members are able to do the right things at the right times. From this perspective, the question of how people are able to get themselves to do the right thing at the right time- how they are able to perform in role- becomes an important and, I suggest, interesting, problem.

There are two aspects to performing in role. One is, in a sense, affirmative- in order to perform in role an individual must know how to do what is required of him. A pianist must know how to play the piece that is on the program for that evening. The other is negative- in order to perform in role, an individual must restrain himself from doing other things of which he is capable but which are irrelevant or counterproductive. The pianist must play the notes that are part of the piece that he is committed to playing, not notes from the many other pieces he knows. In this latter sense, performing in a role always means restricting oneself to expressing only a small part of one's

capacities. People are different from machines in this respect, because machines, generally, have been designed to perform a very narrow range of behaviors and do not have to be restrained, or to restrain themselves in the way that people always do. People, therefore, must possess internal mechanisms that enable them to acquire and express certain of their capacities while, at the same time, suppressing others.

In the discussion below, it will be assumed that a central element in these hypothetical internal mechanisms is the capacity to identify with a given role, that is, to assume that role as an identity, just as an actor assumes an identity in playing a part. The implication of this is that when you take up a role you should be willing and able to become, at least temporarily and in the eyes of others, the kind of person who does that thing and, at least implicitly, nothing and no one else. People never actually do this, I believe, and their identifications with roles can be strong but are never so complete or absolute. As a result, they are always more than just their roles and usually display characteristics that are inconsistent with their roles.

As a result, a tension between the capacities and characteristics of the whole person and the demands of role is intrinsic to organizational life. From one point of

view, this fact may appear to represent an impediment to the smooth functioning of the organization-as-machine because some of the components are unpredictable and unreliable. It is also possible, however, that it is just these extra characteristics that individuals carry into their organizational roles that provide the flexibility and creativity that will enable an organization to adapt to new circumstances and new problems. In any case, the issue is one with important implications for the functioning of any organization, and directs our attention to a psychological substructure that may exert just as much influence on the functioning of the organization as the technological substructure.

The technological substructure, because of its concrete, physical reality, may appear to be easily understood and described. The psychological substructure may be more elusive and difficult to describe. The studies discussed and reported here constitute a test of one possible way of determining at least some of the characteristics of an organization's psychological substructure at a given time.

II. REVIEW OF THE LITERATURE

The studies reported here grew out of the confluence of three areas in the psychological literature. The first is the socio-technical, open systems theory of organizations developed by Eric Trist, A. K. Rice, Eric Miller, and others at the Tavistock Institute. The second is the "computational" or "information-processing" approach to the understanding of perception and mental representation, especially David Marr's work on vision. The third is the psychometric literature on similarity and the analysis of proximity data.

This review will begin with a discussion of the Tavistock tradition, which presents a framework for understanding the problems that confront groups of people who try to cooperate in a common task. The key point, for our purposes, in this section of the review is that a group's ability to perform a task effectively depends, to a significant degree on the ability of the members to identify themselves with the group and with their roles within the group (a psychological phenomenon that the Tavistock theorists called "sentience"). The problem, according to the members of the Tavistock school, is that each individual has the potential for making many different identifications. Some of these may be related to the organization and its

task, but others may reflect the many other commitments that any individual brings into the work situation- the identifications with one's race, ethnicity, or gender being only the most prominent and universal such identifications. Some of these aspects of a person's identity may, under certain circumstances, interfere with their performance in their role, while others may facilitate it. From this perspective, it is important to know what the state of identifications of the individuals in the organization in coming to understand the way in which the organization functions.

It will be suggested here that the claim that a person has identified him/herself with a group is a statement about that person's mental representations and that Tavistock theory therefore ascribes a central role to such representations. Intriguing as this idea is, the theory has little to say about the psychological mechanisms that support and govern the workings of such representations. Their existence and a few of their general properties are inferred from observations of the behavior of individuals in organizations, but little is said about how the specific capacities and limitations of our "mental representational machinery" may shape the ways in which the organizational environment is perceived and behavior in the organization is directed. In this respect, the theory appears to be

incomplete.

The Tavistock theory also suffers from a second, more practical, limitation. In constructing inferences about the relative strength of the multiple possible identifications that may be at work below the surface of organizational life, an observer is forced to proceed in much the way that a person doing psychotherapy does- that is, listening, watching, asking questions, and then using their theories to draw inferences about the underlying situation. Obviously, the method has considerable power, as the successes of the practitioners of such approaches indicate. It will be suggested here, however, that it would be useful to see if it was possible to develop another, independent, method of inferring the underlying mental state of the organization. In some cases, the each method could be used to confirm the inferences or hypotheses drawn from the other. In other cases, it might turn out that each method was capable of revealing a different aspect of organizational life and that a more comprehensive understanding could be attained through the use of both methods. The studies reported here constitute an attempt to propose, develop, and explore such an independent method of determining the identifications of the members of a work group.

Marr's work on vision, the second area covered in this review, provides a remarkably illuminating and straightforward perspective on the problem of mental representation, with direct application to the study of mental representations in organizational life, despite the fact that Marr's work concerned relatively early stages in the processing of information in a single sensory modality. The review will briefly sketch Marr's theory of vision, in order to bring out the general structure of his model, and an attempt will then be made to apply the model to the organizational situation.

We then turn to a discussion of the psychometric literature on the analysis of proximity data. At first glance, this highly technical, emphatically non-clinical literature may seem to have little relevance to the problems that have been discussed here. It will be suggested, however, that the techniques that come out of this research tradition make possible a surprisingly direct form of access to the mental representations that underlay our experience of organizational life.

The "Tavistock" Tradition

The work of Trist, Rice, Miller, and other members of the Tavistock Institute on organizational dynamics has its roots in psychoanalysis and object relations theory, especially the work of Melanie Klein. One of their central concerns was the ways in which social and psychological structures enable us to mediate and manage the interactions between the external environment and our primitive needs. Generally, the characterizations of the "primitive needs" offered by Tavistock school are drawn from Kleinian psychoanalytic theory.

Bion

W.R. Bion was the first of the object relations theorists to do clinical work with groups (Bion, 1961). He suggested that groups operate on two levels:

1) the "sophisticated work group", which is concerned with the explicit (and conscious) task that the group has taken on; and

2) the "basic assumption group", which reflects the underlying, unconscious, assumptions which it (that is, the members, collectively) make about

its task.

Bion offered a classification of the basic assumptions that he had observed in group life, and proposed that the actual behavior of a group should be viewed as a compromise between the demands of both the implicit and explicit tasks.

Socio-technical Theory

Bion began his work in a military hospital and continued it in training groups at the Tavistock Institute. Some of the staff with whom he worked there (including Eric Trist and A.K. Rice) began to apply, and extend, the concepts to their own work in organizational consultation.

They also integrated concepts drawn from open systems theory. As Miller and Rice put it:

> "Any enterprise may be seen as an open system which has characteristics in common with a biological organism. An open system exists, and can only exist, by exchanging materials with its environment. It imports materials, transforms them by means of conversion processes, consumes some of the products of conversion for internal maintenance, and exports the rest. Directly or

> indirectly, it exchanges its outputs for further intakes, including further resources to maintain itself. These import-conversion-export processes are the work the enterprise has to do if it is to live." (Miller and Rice, 1967, p. 3)

This work produced the basic outlines of (and the term) sociotechnical theory. A central point is there are two basic "dimensions" to any organization:

> "... any production system requires both a technological organization- equipment and process layout- and a work organization relating to each other those who carry out the necessary tasks. The technological demands place limits on the type of work organization possible, but a work organization has social and psychological properties that are independent of technology..." (Rice, 1958 quoted in Trist et al., 1963. p. 6)

The degree to which the productive potential of a given technology is realized depends upon the degree to which the sometimes competing demands of the two "dimensions" can be balanced or reconciled:

> "Inherent in the socio-technical approach is the

> notion that the attainment of optimum conditions in any one dimension does not necessarily result in a set of conditions optimum for the system as a whole. If the structures of the various dimensions are not consistent, interference will occur, leading to a state of disequilibrium, so that achievement of the overall goal will to some degree be endangered and in the limit made impossible. The optimization of the whole tends to require a less than optimum state for each separate dimension." (Trist et al., 1963. p. 7.)

This approach is a distinct contrast to the "machine models" of organization that were prevalent at the time (Katz and Kahn, 1966), in which it was imagined that people would fit themselves to the demands of the technical layout.

Some of the first consulting projects of the Tavistock founders began with a similar problem- the introduction of mechanized methods of production in an industry did not produce the increase of production that had been expected. A. K. Rice worked with a group of Indian textile mills, for example, in which a substantial investment had been made in a new type of automatic loom, but "neither productivity as measured in terms of efficiency or quality... was as high as expected... neither was appreciably better than, or always

as good as, on ordinary looms" (Rice, 1958 p. 51). Despite considerable efforts, no technical explanation could be found, and attention was turned to the social-psychological system of the men operating the machines. When changes were made in the social structure which brought the social-psychological system into alignment with the new technical system, productivity increased (Rice, 1958).

Another of the first large scale research projects carried out by the Institute was a study of the English coal mining industry (Trist et al., 1963). This work will be considered here in some detail, as it illustrates some of the basic concepts of sociotechnical theory.

At the time of the research project, the coal mining industry was experiencing a period of considerable technological change. Previously, the prevailing organizational model had been the "single place method", in which each miner was somewhat like a farmer tilling his own plot of land, and performing all of the different functions that are necessary to the enterprise. In the single place method, each miner was responsible for all phases of the work- cutting the coal, carrying it away, advancing the roof supports, etc. When machines such as conveyors and power driven coal cutters were introduced, the work was changed to a larger scale, team operation in which a fairly large group

had to divide up the mining task into its parts and work cooperatively to mine a large coalface (the "longwall method"). Initially, however, the situation was similar to the one that Rice confronted in the Indian textile mills- the increase in production was far less than had been expected, given the obvious superiority of the power driven machinery over manual techniques.

As they examined the situation, Trist and his colleagues came to the conclusion that the difficulty lay in the form of work group organization that was initially used to implement the longwall technology. The technology demanded a high level of cooperation between the individuals responsible for the different aspects of the overall task. On the one hand, in order to get individuals to commit themselves to their particular subtask and to conform their behavior to the requirements of that task, the work organization had to encourage individuals to identify with their particular work roles, defined as:

> "By work roles are meant the jobs which people do every day and with which they become identified- cutters, fillers, pullers, stonemen, etc. In thinking of themselves as such they gradually take on certain common characteristics and may be said to acquire the <u>character</u> of their role. In an

industry such as mining, the role with which a man becomes identified becomes a way of life." (Trist et al., 1963, p.22)

On the other hand, the structure must also encourage individuals who are identified with different tasks to cooperate together- in other words, to see themselves as identified with a common overall task. On the coalface, this might mean, for example, that a cutterman whose work was finished for the moment would "pitch in" to help others remove the cut coal. If he did not, he might end up standing idle for a time, thus reducing the efficiency of the team.

One way to encourage identification with a task is to peg the individual's pay to the performance of that task, and this was characteristic of the early forms of longwall organization. Each group (e.g. cuttermen) was paid according to how much of that task had been accomplished, a measure that encouraged the men to focus almost exclusively on that task. The arrangement did not facilitate cooperation between task groups, however, as there was very little benefit, in terms of wages, to the members of one group for helping out with the work of another:

"Conventional longwall organization is

> characterized by the general conception: <u>one man-one task</u>. Each segregated task group has its own personnel, customs, price agreement, and paynote and is very largely bounded by its own concerns... No overall social organization exists to bind these segregated task groups together in common pursuit of the primary task of completing the cycle." (Trist et al., 1963, p.46)

This lack of cooperation, according to the consultants, explained the disappointingly small increase in productivity for the longwall method. They concluded that the single place method, though it deployed a much less powerful technology, had been supported by a social/psychological system that was so well adapted to that technology that it was fully exploited and workers produced fairly good results. The first longwall methods, which deployed much more powerful tools, were constrained by modes of work organization that were much less effective in exploiting the full power of the technology.

This situation improved very markedly when a new form of organization, the "composite longwall", emerged, in which self-selected groups of about forty men formed which took on full responsibility for operations on a coalface. These groups divided the labor among themselves, and received a

single payment that was determined by the group's productivity. This arrangement, which encouraged a high degree of identification with the overall task of producing coal, rather than with a small part of the overall task, and which encouraged a high degree of cooperation between members of the group, proved to be much more productive, and began to realize the full potential of the power driven technology.

This example illustrates two of the central themes of the sociotechnical tradition. First, productivity is a function of the interrelationship between technology and organization, not of the technology alone. Second, in order to ensure accomplishment of a task, the organizational structure may need, in certain technologies, at least, accomplish two things simultaneously: 1) it must encourage the workers to identify with the particular subtask (e.g. coalcutting) for which they are responsible; and 2) it must encourage a necessary level of cooperation between the different task groups that must work together to accomplish the overall task (maintaining the overall cycle of work on the coalface, in the example above). In a sense, the task structure (the arrangement of roles and tasks in which people are organized to carry out the group's overall task) must be "energized" and supported by an appropriate social/psychological structure if the task is to be

performed efficiently.

"Sentience"

In their seminal work, Systems of Organization (Miller and Rice, 1967), E.J. Miller and A.K. Rice deal explicitly with the problem of balancing the technical and psychological aspects of organizational systems. A key element of their theory of the psychological level (and its relationship to the technical level) is the concept of "sentience", a word chosen by Miller and Rice to denote "the groups with which human beings identify themselves as distinct from task groups, with which they may or may not become identified" (Miller and Rice, 1967). In terms of the argument that is being developed here, "sentience" refers to the degree to which a particular group or category is "energized" in the minds of its members. For example, in the early form of longwall organization described above, task groups such as cuttermen and pullers were accorded a high degree of sentience, while the identification with the overall group on the face was relatively weak. After the pay system was changed, it would appear, the relative sentience of the overall group was greatly enhanced, and the sentience of the sub-groups decreased, and the behavior of the members changed significantly.

According to Miller and Rice, this experience of the identification with a group ("sentience") is fundamental to commitment to task performance and to our ability to mobilize ourselves to perform in a role. In the discussion to follow, such identifications will also be described as "making the group psychologically real."

In Systems of Organization, Miller and Rice examined the ways in which different factors, which could be either intrinsic or extrinsic to the task of the enterprise, affected the distribution of sentience and the performance of the task. They discussed the ways in which different organizational structures and situations create different experiences for the individuals who worked within them, and began to consider the ways in which these different experiences are interpreted and internalized in the mind. They developed their argument by presenting a series of case studies of different types of organizational structure and discussing the different problems that each presents. Two examples are worth considering here.

First, consider the case of the family-owned and operated business in which a group of individuals simultaneously hold, and are strongly identified with, roles in two different systems- the family and the business enterprise. There are some advantages to the arrangement,

such as the high level of task commitment that can be obtained from family members when the boundaries of working life and family life coincide. This is especially important when a business is just getting started and extraordinary efforts are required. There are, however, many difficulties which become much more pronounced as the business grows. For example, as a business grows it may be become increasingly difficult for the family to provide the greatest portion of the labor and capital that is required, but allowing "outsiders" to participate, as investors or employees, can dilute the family's control. Furthermore, the superimposition of two distinct sets of relationships (e.g. a pair who are simultaneously father-son and manager-subordinate) can introduce conflicts between roles that would not exist when the two systems are separate.

The family business is a special case of a more general class of "sentience situations" in which members of an enterprise are also members of other groups which cross-cut the organizational structure and may even extend outside of the organization. Race and gender are the most obvious categories of this type, but there are many others (e.g. political party affiliation). The relative importance of these cross-cutting categories to life within the organization may vary with circumstances, but it would seem likely that they can become very significant at times (race

after a major racial incident, political affiliation in an election year)[1]. It would seem likely that, at least in some cases, these external affiliations could interfere with task performance to such an extent that they would demand special interventions by management.

A second type of "sentience situation" is exemplified by the "interdisciplinary research group", in which individuals who are identified with different disciplines are gathered to work together, for a limited period of time, on a common project. In the beginning, the "sentience problem" may be that the members may all be more strongly identified with the disciplines from which they came than with the new team, and it may be necessary to increase the identification with the team to generate the commitment necessary for the project. Later, however, as the project nears its end, the identification with the group may have become so strong, and the disciplinary ties so atrophied, that the group resists dissolution even though it has, at least in the judgment of presumably objective outsiders, "outlived its usefulness." This sort of problem is becoming

---

[1] It seems likely that the salience of these groupings would be enhanced when membership in the "task-irrelevant" group (race, for example) is correlated with membership in "task-relevant" groups (for example, bosses and secretaries). In such a situation, the identification with the "task-irrelevant" group might accentuate the sense of identification with the "task-relevant" group and intensify inter-group conflict.

much more common in modern organizations which must frequently reorganize to meet the demands of a rapidly changing environment, and in which temporary work groups are not at all uncommon.

Thus, according to this model, each individual in an organization may be subject to the pull of sentience, of identification, from many different directions. He/she will undoubtedly be a member of any number of external groups (race and gender only being the most obvious). Much of the time, perhaps, the external memberships will exert a negligible influence but, at times, the pull may be strong and may complicate performance in role. He/she may also be subject to the pulls of their identifications with multiple internal task groups, as is the case of the members of the interdisciplinary research group. It also seems likely that, over time, events will shift the balance of the multiple forces, tugging the individual more in one direction than another, and that much of the actual behavior of the individuals and of the organization as a whole must be affected by these shifts. Understanding which of the vast array of possible identifications are "psychologically real" at a given moment may be critical to understanding how the organization is functioning at that time.

The Limitations of the Tavistock Approach

The utility of the concept of sentience seems obvious, but there are two limitations that should be considered. First, the initial formulation has not been built upon since the publication of Systems of Organization, and there has been no deeper analysis of the psychology of sentience, and no attempt to understand it in the context of a more general psychological theory. Twenty years later, it appears to be a promising idea that, though it has been used by organizational consultants working in the Tavistock tradition, has not been the focus of a great deal of empirical or theoretical elaboration.

Second, there is the issue of how the distribution of sentience, within a group and across the multiple potential identifications, is determined. Though the question is not directly addressed by Miller and Rice, one assumes that sentience is discerned through the influence that it exerts on the behavior (and, perhaps, reported experience) of the people who inhabit the organization under study. In other words, we watch how people act and what they say about their experience, and infer the existence of sentience identifications in order to make sense of what we see. Even if this method is adequate for practical purposes, it may be insufficiently precise and reliable to serve as a basis for

a development of the concept. Therefore, it might be useful if another, independent, method of determining sentience could be devised, to supplement the inference techniques employed by the Tavistock school of organizational consultation. To the extent that the two methods produced the same findings, in a given situation, they would provide confirmation for each other. It might also be possible that each method would have some special sorts of information that only it could reveal, and that the two methods might illuminate some different aspects of organizational life. In a way, what is being suggested is something that is similar to the relationship between a clinical interview and psychological testing- each method is directed toward the same goal, and one expects the findings of the methods to be consistent with each other, but each method is also capable of bringing to light aspects of mental functioning that are inaccessible to the other.

Ideally, such a method would attempt to directly investigate the groups into which people (mentally) place themselves and others in their organization (in other words, to attempt to develop a "measure" of sentience), and explore the correspondences between the measurements taken and: 1) the formal task structure; and 2) external events that might be expected to have an impact on group identifications. Something of that sort was attempted in the studies reported

here. The method has the significant advantage that it makes it possible to apply a very powerful statistical technology, based on the concept of similarity, to these issues. Before doing that, however, it will be necessary to become a bit clearer about mental representations and their relationship to behavior and experience than we have been in this discussion so far.

A Computational Perspective on Mental Representation

David Marr was a neurobiologist and cognitive scientist who spent most of his career at the Massachusetts Institute of Technology. Marr's central concern was the nature of vision, and his goal was to develop a model of visual processing that was sufficiently precise and detailed that it could be modeled on a computer. Perhaps the best starting point for an understanding of his work is a respect for, even a sense of wonder at, the remarkable a feat that is performed by our visual sense. This "sense of wonder" is not so easily attained, probably just because vision is so rapid and effortless a process for us that it appears "easy." It may be useful, therefore, to stop and reexamine the task that is accomplished by the visual system.

The Problem of Vision

Imagine that we are sitting in a room. The room is, in fact, lavishly furnished- the walls are cluttered with pictures, there is a richly textured Oriental rug on the floor, and the room is full of ornately carved wooden furniture upholstered in a variety of patterned materials. It is, however, pitch black. There is no light at all and, as result, we can see absolutely nothing. Suddenly, a single candle is lit.

As the candle burns, light is emitted by the flame, spreading out in all directions. The light spills out over the room, encountering the surfaces of all of the many different kinds of objects that are to be found there. In the ensuing interaction between the surface and the light, the light is transformed. Some parts of the light are absorbed (depending on the nature of the surface) while others are reflected (depending on the nature of the surface). In a sense, the light, as it spreads through the room, is "bruised" by its contacts with the surfaces and, in this "bruising," it becomes imbued with information about the surfaces of the objects.

Eventually, a rather small part of the candle light reaches our eyes and, after being focused by the lens, it

strikes the retina, where it encounters a layer of photoreceptive rod and cone cells that constitute the outermost layer of the nervous system. The behavior of these cells is greatly influenced by the character of the light that strikes them, and their behavior exerts an influence, in turn, on the other cells in the nervous system.

Less than half a second later, we see the room, and we see it in what might be called incredible detail, were all of this not such a common occurrence. We see a space that is filled with discrete objects and we can see where all of the objects are, in three dimensional space- where they are in relation to our body, roughly how far away they are, which ones are in front of others, etc. We can see the fuzziness of the rug, and the hard, smooth shine of a varnished tabletop. The patches of oil paint on the pictures are seen as images of people and landscapes. In a remarkably brief time, the visual system has taken the "bruised candle light" that was collected by the retina, extracted the information that was impressed on the light in its interactions with the objects in the room, and used it to construct a vivid, highly detailed picture of the room and its contents. Furthermore, if the candle is suddenly blown out, and the room is dark again, we can still "picture the room in our imagination" in considerable detail. In

other words, we still have a mental representation of the room, based on the visual image, available.

For Marr, the problem of vision is an information processing problem- how can the nervous system construct such a detailed representation of the room from the pattern of light intensities that strikes the retina? He took pains, however, to emphasize that internal representations are an integral part of an information processing account:

> "Vision is the process of discovering from images what is present in the world, and where it is... if we are capable of knowing what is where in the world, our brains must somehow be capable of representing this information- in all its profusion of color and form, beauty, motion, and detail. The study of vision must therefore include not only the study of how to extract from images the various aspects of the world that are useful to us, but also an inquiry into the nature of the internal representations by which we capture this information and thus make it available as a basis for decisions about our thoughts and actions." (Marr, 1982, p. 3)

Interestingly, the problem that Marr was addressing can be put into the "open systems" terminology discussed in the previous section. The visual system has, as its input, the light-impregnated-with-information that is collected at the retina and, as its output, our visual experience of the three-dimensional space that surrounds us, segregated into objects. From this point of view, the task that Marr set himself was to describe a set of conversion processes that could transform the input into the output.

A Modular Model of the Mind

Marr's solution was to propose a series of interconnected processing modules, each of which performs one part of the task of the visual system:

1. the "Image" in which represents the intensity of the light that falls in each part of the retina;

2. the "Primal Sketch" in which information about the pattern of light intensities is collated;

3. the "2 1/2 Dimensional Sketch" which contains information about the depth and orientation of visual surfaces; and

4. the "3-D Model Representation" which provides information about the shape of objects, their location and orientation, and the volume which they occupy.

The inputs for each module are provided by the one that precedes it in the sequence, and its output constitutes the input for the next module.

The details of Marr's theory of vision need not concern us, but there are three points about his general approach, and the view of the mind that it reflects, that will be useful for our purposes.

First, if we accept the notion that mental processing systems can be regarded as systems of interconnected information processing modules, it would seem reasonable that the approach could be extended beyond the visual system to higher levels of cognitive processing. Figure 1 presents a schematic diagram of such a system. The primary inputs for the information processing system are the various forms of sensory stimulation (light, sound, smell, etc.) that have been imbued with information about the external world. These inputs are registered by sensory organs such as the retina, in which the extraction of information is begun. This information is then passed on to other modules for

Figure 1:
A Schematic Diagram of a Modular Information Processing System

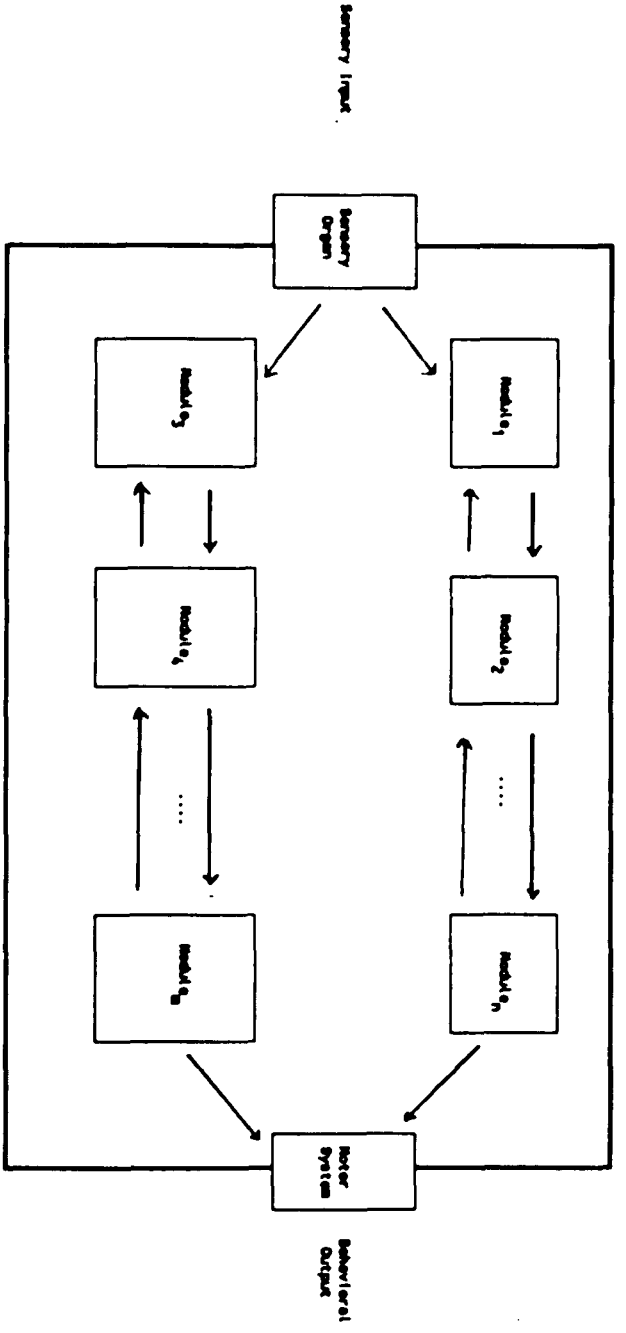

further processing. As indicated in the schematic, some modules may receive input from, and send output to, more than one module. One likely instance of this would be where information from several sensory modalities is combined. Furthermore, there would seem to be no a priori reason why information cannot be sent to modules that occupied earlier places in the processing sequence as well as later ones (attention, for example, might involve a "tuning" of earlier stages by later ones). Ultimately, output from the system of modules provides the basis for experience and behavior, which constitute the final outputs of the total system.

The essential feature of Marr's general approach is that the total information processing system can be analyzed into a number of sub-systems. Within this general framework, he proposed a highly specific architecture of modules for the early stages of visual processing, with details about the inputs, outputs, and processing mechanisms of each module. When we move beyond vision, however, Marr provides only general guidance, and the number and nature of, and the interrelationships between, the modules in the rest of the total information processing system are unknown.

The second point is that there is no reason to suppose that all stages of processing are accessible to conscious inspection or that they all have direct behavioral

manifestations. Some are simply intermediate stages of processing that provide information that later becomes the basis for behavior or other representations that are accessible to consciousness. This is clear in the case of vision- some of the early stages in the system (the Image, for example) are not a part of conscious experience, but they do provide input for later stages that are. In fact, it may well be that it is only a few of the very late stages of cognitive processing that form the basis of behavior and experience. This fact has an important implication. If experience and behavior are, most likely, greatly determined by the architecture of the earlier stages of processing, then it will be necessary to understand that architecture in order to fully understand them.

One way to do this is to do what Marr did, that is, to work "downstream" from the initial inputs. Another is to find ways of working backwards ("upstream") from the final outputs (experience and behavior) to the intermediate modules that precede them. In general, one might suppose that working downstream will be more useful in understanding early stages of processing and that working upstream will be more useful in understanding later stages, which, in turn, suggests that an upstream approach would prove most fruitful for understanding our representations of organizational situations.

The third point is that each of stage in the process can be treated as a function that maps a particular input onto a particular output. From this point of view, the system can be seen as a series of interlocking functions, through which an initial input is passed, where each stage transforms the throughput in a characteristic way.

This suggests a strategy for the sort of working upstream described above. First, we identify a behavioral or experiential output that is easily accessible and which taps the (inaccessible) internal representation that is of interest to us. Then, we formulate a hypothesis about the nature of the function that maps the internal representation onto the behavioral or experiential output. Finally, we construct an inverse of that function and apply it to the behavioral/experiential output, in the hope that what will emerge is a reconstruction of the internal representation. (This is, in fact, exactly what we do when, in our clinical work, we make inferences about internal states on the basis of observations of behavior and reports of experience.)

In the section that follows, we will examine a candidate for such a function.

The Psychometrics of Similarity

The concept of similarity has played a central role in many psychological theories, though it has not always been explicitly acknowledged. Consider, for example, the general form of a behaviorist account of learning:

1) a subject is exposed to repeated simultaneous presentations of two stimuli (e.g. a bell and the dispensing of food);

2) an association forms between the two stimuli; so that

3) when one part of the associated pair (the bell, perhaps) is encountered at a later time, the other part (the food) is recalled to mind and an appropriate response (salivation) is evoked.

Explanations of this form, which dominated American psychology for many years, turn on the concept of similarity. How does the subject determine that the bell encountered at a later time is the same as the earlier one, which was involved in the association? Bells may differ in a number of respects, such as tone, loudness, or duration, and the chances are that the second bell will not be

identical to the bell (i.e. absolutely the same in all respects) that was involved in the creation of the association. Furthermore, unless the model allows for the response to be evoked by bells that are at least a "little bit different" from the original bell, it will explain such a narrow range of phenomena that it will appear to be useless. On the other hand, if the model does not require the later bells to be "very much like" the original bell, so that stimuli that are "very different" from the original bell are allowed to evoke the response, it becomes difficult to see why the original association was important and the point of the explanation is lost. This has been called the "stimulus generalization problem" (see Shepard, 1957; Shepard, 1958a; Shepard, 1958b for a discussion). In order for explanations of this form to be useful, the association extend itself to a certain degree, but not too much.

Obviously, the solution must be found in the notion of "similarity." Stimuli that are perceived to be "similar" to the original bell should be expected to evoke the response, and stimuli that are perceived to be "dissimilar" should not. If this solves the "stimulus generalization problem," however, it opens a new series of questions about the nature of similarity. For reasons such as this, similarity has been the focus of a long tradition of research in American psychology.

Attneave and the Origin of the "Spatial Model"

Little systematic work was done however, until Fred Attneave's 1950 article on similarity. Attneave began with the observation that "it is obvious that when things are similar they are similar with respect to something." In other words, the similarity between two stimuli is a function of certain qualities that they possess, and it seems reasonable to suppose that our concept of similarity will be shaped by the way in which we conceptualize these qualities. The "something" may, he says, "be conceptualized as either more or less discrete and common elements or as dimensions on which the objects have some degree of proximity." Attneave was the first theorist to distinguish between the two different approaches to similarity that have dominated the subsequent literature.

The first is the "categorical" approach, in which the qualities are thought to be certain common or distinctive features of the stimuli that can be thought of independent categories. For example, in judging the similarity of two people, we might note that they both had red hair and green eyes (and were similar in those respects) but that one was tall and female while the other was short and male. Our overall judgment of the similarity between the two would

depend on the way in which we combined and balanced these different qualities.

The second is the "spatial model", in which each quality is treated as a continuous dimension. For example, consider the bells that were mentioned above. Each bell has a certain tone (i.e. frequency), a certain loudness (measured in decibels), and a certain duration (measured in seconds). Each of these three qualities can be thought of as a scale of measurement, as a ruler is a scale measuring length, and we can use them to "take the measurements" of any given bell. The similarity between two bells with respect to a single quality (e.g. loudness) is determined by how similar their measurements on the loudness scale are, or the distance between their locations on the scale. If the "rulers" are thought of as the axes of a coordinate space, we can assign each bell to the specific point in the three-dimensional space that corresponds to its measurements.

Any theory of similarity must also deal with what we might call the "combination problem"- an overall similarity rating is obtained for two stimuli that differ with respect to more than one quality. In order to make a judgment of "overall similarity" subjects must take into account several types of similarity and difference information and somehow "add them up together." One of the appeals of the spatial

approach is that it offers a simple, intuitively appealing model for how this could be done. As discussed above, stimuli are conveniently treated as occupying particular locations in a multidimensional space. The difference between stimuli can thought of as corresponding to the length of the line connecting the two points in that space[2].

Without extensive argument, Attneave elected to pursue the "spatial" approach[3]. His experiments dealt with artificially constructed, relatively simple perceptual stimuli (e.g. shapes, colors) that were quite suited to the spatial model, and his results supported the usefulness of that model.

For our purposes, the significance of Attneave's work is that it identifies the two fundamental questions that any theory of similarity must address:

---

[2] In his discussion, Attneave refers to the work of Richardson and of Young and Householder, in this area. Young and Householder (1938) developed the first procedure for constructing a spatial configuration of points from an error-free matrix of inter-point distances, using a technique similar to factor analysis.

[3] The other possibility, the road that Attneave and most of the other early workers in the field did not take, corresponds to the "set theoretic" approach later developed by Tversky, Hartigan, and others (see below).

1. The spatial versus feature issue- that is, are the qualities that are taken into account in estimating similarity to be treated as continuous dimensions or as discrete categories?; and

2. The "combination problem"- that is, how are the different qualities that the stimuli possess to be combined together when a judgment of overall similarity is made?

The different statistical techniques that will be discussed below (multidimensional scaling, hierarchical clustering, additive trees, additive clustering, etc.) can be thought of as procedures for fitting similarity data to different models, each determined by the choice of a specific set of answers to these two questions.

Reconstructing Perceptual Structures from Similarity

In this discussion, our attention has been focussed on the ways in which relatively well-understood qualities (e.g. the different aspects of the sound of a bell) can determine perceived similarity. Attneave, for example, constructed artificial stimuli, in which he could control the perceptual qualities possessed by the stimuli, and studied the impact that various changes had on perceived similarity. With most

naturally occurring stimuli, however, one cannot know in advance what the key perceptual features are, and this may, in fact, be just what one wants to discover. Much of the considerable interest in similarity that was shown by the community of psychological researchers was due to the possibility that one could use it to discover hitherto unknown psychological structures that govern perception and categorization. Suppose, for example, that one wants to know how people perceive the sounds used in speech. One approach would be to work backwards from judgments of perceived similarity, which, in many cases, are easily obtained. What was needed, then, was a technique for reconstructing the structures that are responsible for the perceptions of similarity from the ratings themselves. The interest in such techniques stimulated a great deal of research.

The logic that underlay this effort is central enough to the rationale for the research reported here that the following digression, which is intended to illustrate the problem and its solution, may be justified.

This problem is, in a sense, comparable to the following, more mundane, problem. Imagine that we have an ordinary highway road map of the United States. On the map are many cities (Albuquerque, Chicago, New York, Los

Angeles, etc.) spread out across the map, each in a different location, each connected to the others by a system of roads. We use the map for number of purposes: to choose routes, estimate distances and times it will take to travel between cities, etc.

On many maps, we will also find a "mileage chart", a table that gives the distances between all of the different cities, to be used in estimating travel time. The table is organized as a triangular matrix. To find the distance between two cities, we read down the column for one city until we come to the row for the other city, and note the number that we find at the intersection[4].

We presume that this mileage chart was constructed by someone who used the map to measure out the distances needed to fill out the chart. So, we imagine, the distances in the chart are derived from the representation that is the map. But suppose that the map was inaccessible to us but that the mileage chart was still available. Could we reconstruct the map from the chart?

---

[4] The matrix need be only a triangular one because we know that the distance from A to B is the same as the distance from B to A, in this situation, at least. We can, however, imagine situations, places where there are a lot of one-way roads, for example, in which this would not be true and we would need to fill in the whole chart.

The answer is that we can, at least as long as know, or can assume, certain facts about the original map. For example, we need to know what the numbers on the mileage chart mean. (In this case they represent spatial distance, but we will not want to limit ourselves to such instances. We should be open to the possibility that they mean something else, such as the difference in the populations of the two cities, for example.)

Even if we know that the numbers express spatial distance, we have to know the dimension of the space. Is it a one-dimensional space, such as distance along a railroad line? Or a three-dimensional space, such as the distance between offices in a building?

A third thing that we may need to know is how distance is determined. Is it the length of a straight line connecting the two points in question (as in a chart that gives the airline distances between cities) or is it the length of the route that one actually has to travel (as may be on a highway map)?

If we know all of these facts about the original map, we can partially reconstruct the map. One method, for example, would be to:

1. Begin by choosing any two of the cities (A and B) and drawing a line between them that represented the distance as indicated on the mileage chart.

2. Then, take any other city (C) and draw circles, centered on A and B whose radii are the distance from that city and C.

3. If those circles intersect at only one point, C is located there. If they have two points of intersection, choose any one, and place C there.

4. The location of the rest of the cities can be found by drawing the circles centered on the cities whose location has been determined that correspond to the distances between those cities and the city whose location is unknown.

This method is a very simple one and depends on certain critical assumptions. For example, it depends on the assumption that the distances on the mileage chart are perfectly accurate, which is, of course, not always the case. On the other hand, the reader may have noticed that the complete mileage chart contains much more information than is needed to locate all of the cities using this method. In instances where one cannot assume that the

estimates are this accurate, one has to use the redundancy in the data to correct for the error in the distances.

The method described above is a very simple example of a procedure for fitting distance data to a specified type of map in order to reconstruct the particular map (of that type) that could have generated the distance data. The statistical procedures that will be discussed below can all be thought of as methods for fitting the (errorful) distance data to different types of "map." It is also worth noting that these procedures can be thought of as candidates for an inverse of the "similarity function." When we make similarity judgments, we draw upon some sort of (usually unknown) internal representation, which provides the input for the process of arriving at a similarity judgment, and the judgment itself is the (accessible) result of that process. In other words, similarity can be thought of as a function that maps data from an internal representation onto a scale of perceived similarity. Procedures such as that described above (and the others that will be discussed below) can be thought of as inverses of particular versions of the similarity function. In this way, we can work "upstream" from an observable behavior to the unknown internal representation.

Multidimensional Scaling (MDS)

Attneave was followed by many investigators who pursued the spatial approach to the study of similarity. The response was somewhat delayed, however, by the lack of a convenient procedure for reconstructing the spatial arrangement from the similarity data. This situation began to change when Torgerson developed the first "routine" procedure (Torgerson, 1952) which, furthermore, could be used with errorful matrices of interval-scale estimates of distance. However, the method imposed strict requirements on the data- "that the perceptual or cognitive structure of the set of stimuli be Euclidean in nature... [and] that observations on similarity of pairs of stimuli be linearly related to distances between points in the space" (Torgerson, 1965). These limitations, and the difficulty of the procedure itself, limited the extent of its applications.

The procedure became much more convenient after the publication of Shepard's method of non-metric multidimensional scaling (Shepard, 1962a; 1962b), which, as Torgerson said, "provided the first computerized, iterative program" (Torgerson, 1965) for reconstructing the spatial arrangement of points, while imposing significantly less

stringent requirements on the data[5].

A further improvement was contributed by Kruskal (1964), who developed an alternative algorithm (using an "optimization" method) with some distinct computational advantages. He also developed a measure called "stress" which provided another indicator of the "goodness-of-fit" of a given MDS solution.

Individual Differences Approaches to MDS

The next major extension of the MDS model was the development, by Carroll and Chang in 1970, of a workable generalization to an individual differences model. In part, this effort was an attempt to deal with a technical problem confronted by the MDS technique.

Typically, the proximity data (the estimates of the psychological distance between stimuli) used in the procedure were gathered from a number of subjects. Each subject might be asked, for example, to rate the similarity between all possible pairs of a set of objects, and these ratings are then arranged into the triangular matrices (one

---

[5] Like Attneave, Shepard was interested in multidimensional scaling because of its relationship to the "stimulus generalization problem."

for each subject) that are the most frequent form of input for MDS. The investigator must then decide how to use the matrices for the individual subjects. There are two options.

First, each individual matrix can be submitted to the MDS procedure, producing a spatial arrangement of the stimuli that accounts for the proximity data. The investigator will then have a spatial arrangement for each subject, and will have to inspect them to see if they are consistent, which is exceedingly difficult. Furthermore, if the arrangements are inconsistent, it is not clear how this can be interpreted. The second option is to combine the individual matrices (most often by averaging the individual distance estimates to produce a single estimate), create a single matrix, and then use this in the MDS procedure. This second option is simple, but it is open to question. Averaging is only valid if the subjects are using representations that are basically the same in three respects: 1) they using the same set of characteristics as "dimensions"; 2) they assign the stimuli to the same places on the dimensions; and 3) they employ the same "rules" for combining the different dimensions into an overall estimate of distance. As is generally the case with the MDS model, the assumptions seem plausible when the stimuli are thought to be simple, as in psychophysics, but appear more doubtful

when complex, conceptual stimuli are used.

The first attempt to deal with this problem was advanced by Tucker and Messick (1963). In their procedure, called "Points of View", each subject's ratings are correlated with every other subject's, and the resulting correlations are used to create a correlation matrix which can then be subjected to factor analysis. In a sense, the correlations are distance estimates (they reflect the degree to which subjects rate the stimuli similarly), and the technique attempts to identify clusters of subjects who employ the same "strategy" in making their distance estimates (the "Points-of-View"), and to then use MDS to obtain a spatial arrangement for each cluster. The technique was cumbersome and was not widely used.

Though other approaches were suggested in the next few years (McGee, 1968; Horan 1969), the INDSCAL method of Carroll and Chang (1970) has been by far the most widely accepted. In their model, it is assumed that all of the subjects use the same set of dimensions, and that the relative placement of the stimuli along the dimensions is also the same. Subjects differ in the relative weights that they give to each dimension when combining them into an overall estimate of distance. One might visualize this by imagining that all subjects are using the same spatial

arrangement, but that each subject "stretches" the space, along the axes of the dimensions, in their own way, and that the "stretching" accounts for the differences in the ratings. Thus, the output from the INDSCAL method is a MDS-type output, assigning each of the stimuli to a particular location in a multi-dimensional space, and, for each subject, a set of weights indicating how much influence each dimension has in determining that subject's overall estimates of distance. The INDSCAL approach provided a conceptually (and computationally) convenient solution to the problem of individual differences.

As a result of these technical advances, the MDS model became easy to apply, and the next two decades saw a considerable "boom" in applications of the MDS technique to a wide range of data (see Torgerson, 1965; Shepard, 1974; Carroll, 1976; Kruskal and Wish, 1978 for reviews).

As the MDS technique came into general use, there was a great deal of discussion of the strengths and weaknesses of the spatial approach, both by its proponents and critics. For the most part, discussions by the developers of the technique focussed on technical problems and on the limits of the then available techniques for fitting data to the models (Torgerson, 1965; Shepard, 1974).

Non-Spatial Theories of Similarity

At the same time, however, some theorists were identifying and questioning the assumptions that underlay the spatial approach and developing alternative, non-spatial models of similarity and statistical approaches that they believed to be superior to multi-dimensional scaling. For the most part, these investigators were exploring the "categorical alternative" that Attneave had chosen to ignore, that is, conceiving of the different aspects of similarity as independent categories, rather than as orthogonal spatial dimensions. In general, these alternative approaches are characterized by the use of "set-theoretic", rather than "spatial", models, and lead to the development of "clustering", rather than multidimensional scaling, statistical procedures (Sokal and Sneath, 1963; Johnson, 1967).

The distinguishing characteristic of the set-theoretic models is that they assume that the differences (and similarities) between stimuli arise because the stimuli possess different discrete features or, in a manner of speaking, they belong to different sets or categories. As an illustration of one "clustering" approach, consider a

group of men who are all members of the same club[6]. This membership creates a degree of commonality among them, a degree to which they are all similar. In addition to this membership, however, each individual is a member of a very large number of other groups. Some may be Republicans, some Democrats, and some may have no political affiliation. They will each be members of different professions. Some will be husbands, some single. From the clustering point of view, the degree of overall similarity between two individuals is a function of: 1) the number of common memberships they have, which increases similarity and decreases psychological distance; and 2) the number of different memberships they have, which increases psychological distance and decreases similarity.

Though there are ways in which MDS and clustering approaches can be shown to be fundamentally similar ("dimensions" are certainly similar to "categories", for example) they are usually quite different for practical purposes. For one thing, the MDS approach is inclined to regard differences between stimuli as having arisen from quantitative differences on a small number of (qualitative)

---

[6] This is an example of a specific type of clustering structure, an "additive clustering" structure in which an individual object can be a member of multiple, cross-cutting categories. There are other, more restrictive, types of clustering structures (e.g. hierarchical clustering structures).

criteria. In other words, they are inclined to exploit differences in placement along the dimensions to account for a lot of variance, and to keep the numbers of dimensions very low. Clustering models, on the other hand, are inclined to try to account for similarity variance by hypothesizing a large number of qualitative distinctions (the categories) and place much less emphasis on quantitative distinctions within a category.

Although some work was done, in the mid-1960's, on statistical methods for fitting data to clustering models, the early techniques did not attain the popularity that MDS methods did. In the 1970's, however, there was an increasing realization that the clustering approach might well be better suited to certain types of data. The MDS developers themselves began to speculate about the usefulness of combining clustering and spatial techniques into so-called "hybrid models" to more accurately reflect the latent structure of many interesting sorts of distance data (Carroll, 1976).

Tversky's Contrast Model

The most thorough critique of the spatial approach was advanced by Beals, Krantz, and Tversky. In a series of papers written in the late 60's (Beals and Krantz, 1967;

Beals, Krantz, and Tversky, 1968; Tversky and Krantz, 1970), they undertook a rigorous critical examination of the axiomatic foundations of the spatial approach that led them to question its applicability in many of the situations of interest to psychologists. In 1977, Tversky published (Tversky, 1977) a major paper[7] in which he proposed a non-spatial alternative theory of similarity judgments, called the "contrast model."

Tversky begins his exposition of the contrast model with a brief discussion of the assumptions that underlie the spatial approach. He argues that:

> "... most theoretical and empirical analyses of similarity assume that objects can be adequately represented as points in some coordinate space and that dissimilarity behaves like a metric distance function. Both dimensional and metric assumptions are open to question." (Tversky, 1977, p. 327-328.)

With respect to the dimensional assumption, Tversky simply asserts that "it seems more appropriate to represent

---

[7] Tversky titled his paper "Features of Similarity", presumably an allusion to Attneave's "Dimensions of Similarity."

faces, countries, or personalities in terms of many qualitative features than in terms of a few quantitative dimensions." (Tversky, 1977, p. 328) His discussion of the metric assumption is more detailed. He cites three axioms that any metric distance function must satisfy[8] and identifies circumstances in which each of them, he argues, can be shown to fail to hold. Tversky then develops his alternative model and shows how it can account for the situations in which the metric axioms fail.

The assumptions that underlie the contrast model are quite different from those that underlie the spatial/dimensional approach. First, the contrast model assumes that the objects that are being compared (in a similarity rating task, for example) can be represented as sets of distinct characteristics or features, which "may correspond to components such as eyes or mouth; they may

---

[8] The three axioms are:

i. minimality-

$d[a,b] \geq d[a,a] = 0$ that is, that the distance between an object and itself is always zero;

ii. symmetry-

$d[a,b] = d[b,a]$; and

iii. the triangle inequality-

$d[a,b] + d[b,c] \geq d[a,c]$

represent concrete properties such as size or color; and they may reflect abstract attributes such as quality or complexity" (Tversky, 1977, p. 329). These features may be binary, implying that an object either has or does not have possess the characteristic (in which case possession of a feature can be thought as corresponding to membership in a category), but they may also be qualities that vary over a continuum. Because a very large number (an infinite number, for all practical purposes) of characteristics can be attributed to any given object, Tversky says, "when faced with a particular task... we extract and compile from our data base a limited list of relevant features on the basis of which we perform the required task" (Tversky, 1977, p. 329).

Thus, in judging the similarity of two objects, the subject must first select a relatively small set of features to characterize the objects. With any two objects, A and B, there will be two different types of features that can be used in the judgment: 1) features that are possessed by both A and B ("common features"); and 2) features possessed by A and not by B or by B and not by A ("distinctive features"). Intuitively, it would seem that the greater the number of common features, and/or the greater the salience of those features, the greater the similarity. By the same token, the greater the number of distinctive features, and/or the

greater their salience, the greater the degree of dissimilarity.

Tversky's formula for similarity, following this line of thought, treats it as a function of these three types of features:

$$S(a,b)=\theta f(A \text{ and } B)-\alpha f(A-B)-\beta f(B-A)$$

where:

a and b are the objects being compared;

S(a,b) is the similarity between a and b

(A and B) is the set of features common to both a and b;

(A-B) is the set of features unique to A;

(B-A) is the set of features unique to B;

α,β,and θ are weights reflecting the importance of each of the three types of features; and

f is a scale reflecting the salience of the

different features.

Statistical Techniques and the Contrast Model

Tversky also demonstrated that all clustering models can be shown to be special cases of the contrast model, obtained by varying the parameters of the model. Two of these variants, additive clustering and additive trees, are especially relevant to the problem under consideration in this research.

To derive the additive clustering model, we hypothesize that persons performing the similarity task rely only on common features, and disregard distinctive features (in the terms of the model, that $\alpha=\beta=0$). In this case, similarity is the sum of the common features, each multiplied by a weight reflecting its salience. If we imagine that each feature corresponds to a category, then the similarity between two objects (a and b) is the is the sum (weighted to reflect salience) of all the categories to which they both belong. Objects that are both members of highly salient categories will be highly similar, and objects that have fewer, less salient common "memberships" will be less similar. Shepard and Arabie (1979) developed an algorithm (ADCLUS) for fitting similarity data to this additive clustering model. Subsequently, Arabie and Carroll (1980)

devised an alternative algorithm (MAPCLUS) with some computational advantages, and Carroll and Arabie (1983) generalized the additive clustering model to take individual differences in account (INDCLUS), using the same sort of approach that had been used in developing INDSCAL.

If we assume that only distinctive features are taken into account in making similarity judgments (i.e. that $\theta=0$), the contrast model reduces to an additive tree model. Sattath and Tversky (1977) developed an algorithm (ADDTREE) for fitting similarity data to this model, and Corter and Tversky (1987) developed a variant (EXTREE) in which additional clusters that cross-cut the nested structure of the tree can be superimposed on an additive tree structure.

A Preliminary Summary and Integration of the Literature

The research on similarity has produced a rich array of techniques for fitting to proximity data to a specified model. Furthermore, these techniques have the practical advantage that the investigator is not required to choose a model before collecting the data, as the same data can be analyzed by all of the different techniques, and the different solutions can be compared in terms of their interpretability and their goodness-of-fit.

At this point, it may be appropriate to attempt to sketch out the possible relationships between the Tavistock theory of organizations and this psychometric tradition.

In brief, the argument consists of the following steps:

1. The accomplishment of large-scale tasks requires the cooperation (the coordinated activities) of multiple individuals. In order to facilitate this, the whole task is divided into a number of inter-related subtasks. This division of labor is what we refer to as the "task structure."

2. The task structure defines roles, which are the functions that individuals need to perform.

3. In order to get themselves to perform the task, to perform in role, individuals must have some way to retrieve or recreate a self-representation that guides their actions.

4. In so doing, they become the role- they identify with it, assume it as an identity.

5. They also accept a set of representations of other

individuals, who are performing in their complementary roles, and who take on, in one's mind, identities that are determined by those roles.

6. We hypothesize that there exists a mental representation whose function is to represent the self and others in the organization in relationship to each other. This "mental map" of the organization contains all of the information about individuals and all of their perceived identifications (including both task-relevant identifications and task-irrelevant identifications).

7. When asked to judge the similarity of a number of individuals drawn from the same group, subjects draw upon this "mental map." Their first step is to select a relatively small number of characteristics (out of the virtually infinite number of possible characteristics) to be used as the determinants of similarity. Judged similarity is a weighted function of these characteristics.

8. The statistical techniques that have been developed to reconstruct representations from

> similarity judgments may enable us to reconstruct the identifications that are contained in the mental representations that people have of the organizations they inhabit.

Results from previous research and the two studies to be reported here provide considerable support for this line of reasoning.

Similarity and the Mental Representation of a Social Environment

There have been very few studies that applied the statistical techniques discussed above to proximity data drawn from social environments. One study, by Breiger, Boorman, and Arabie (1975), used data on observed social interaction in a group of 14 workers to construct a proximity matrix, but this approach is fundamentally different from the approach described below in that it relies on observed behavior rather than on the mental representations that are presumed to be the basis for similarity judgments. A review of the literature revealed only one study, carried out by Jones and Young at the Thurstone Laboratory, that used the statistical analysis of similarity ratings to uncover aspects of social structure

(Jones and Young, 1971, 1972). This study is worth considering in some detail, as the two studies reported here are modeled so closely on it.

Jones and Young were interested in: 1) the attributes that individuals use in perceiving the other inhabitants of a common social environment; 2) the relationship between differences in the roles of the subjects and the different weighting of the different attributes in making judgments; 3) the ability of the configuration revealed by the multidimensional scaling (INDSCAL) analysis in predict the behavior of the group members; and 4) the stability of the social structure revealed by the scaling analysis over time.

They argued that the MDS-type approach to the determination of social structure had several advantages over traditional sociometric methods: 1) the subject chooses the "aspects of the stimulus person he will judge", rather than having the investigator arbitrarily impose aspects; 2) "the interpretation of the dimensions is objective" because the interpretations of the dimensions can be tested by correlation with other rating scales; and 3) the INDSCAL model allows for individual differences between the subjects to be included in the model.

The subjects were the students, faculty, and staff of

the L.L. Thurstone Psychometric Laboratory of the University of North Carolina. The group of raters was essentially identical with the group of stimulus persons. Subjects were asked to: 1) rate their own degree of familiarity with each of the stimulus persons; 2) judge (on a nine point scale) the similarity of all possible pairs of stimulus persons; 3) rate each of the stimulus persons on several unidimensional, nine-point scales (e.g. status, perceived degree of interest in various professional activities) which constituted the "hypothesis as to the identities of the dimensions of interpersonal perception"; and 4) select the individuals from amongst the stimulus persons that they associated with most and least frequently and that they judged to be most and least similar to themselves. The experiment was conducted twice: in February 1969 and February 1970.

The data were analyzed to: 1) determine the group subject and stimulus spaces for each of the years; 2) test the interpretations of the dimensions of the social structure revealed by correlating them with unidimensional rating scales; 3) search for "systematic differences in interpersonal perception as a function of role and status differences among judges"; and 4) determine "the relationship between interpersonal perception and interpersonal behavior."

A three-dimensional solution was obtained. Based on correlations with the unidimensional scales, the dimensions were interpreted as: 1) "status"; 2) "political persuasion" (left-right, liberal-conservative); and 3) "professional interests" (content vs. methodological/statistical). The authors explain the emergence of these dimensions:

> "Effective functioning in this group demands that its members be reciprocally aware of another's research and teaching interests. Similarly, awareness of the status hierarchy can be assumed to be intrinsic to the functioning of this or any other group where formal or informal differences in status exist. Likewise, the emergence of the political persuasion dimension was not surprising. The study was conducted during a period of campus and national unrest..." (Jones and Young, 1972, p. 119-120)

Analysis of individual differences revealed significant differences between faculty, senior graduate students, and junior graduate students in their weighting of the different dimensions in making their similarity judgments: faculty were "more status conscious"; senior graduate students "weighed professional interest more heavily"; and "new graduate students and clerical personnel paid more attention

to political persuasion." Interpersonal behavior could also be correlated to dimensions of the INDSCAL solution.

These results are interesting for two reasons. First, they demonstrate that the analysis of similarity data, through techniques such as multidimensional scaling, can reveal meaningful aspects of social structure, and that the structure revealed can be interpreted in terms of task and sentience structure. Second, they support the idea that task-irrelevant categorizations (in this case, political persuasion) can intrude to a detectable degree into the perception of psychological distance. This, in turn, supports the hypothesis that this approach might have some diagnostic effectiveness.

III. THE PILOT STUDY: THE ADMINISTRATION OF AN ACADEMIC HEALTH CENTER

As a first step in investigating the issues raised in the above review, a pilot study was also conducted, using a design similar to that used by Jones and Young. The study had two purposes. First, it was designed to confirm the findings of Jones and Young about the usefulness of the analysis of similarity data in reconstructing the mental representations of social space that are held by members of a group. Second, it provided an opportunity to fit the data to models (e.g. additive tree, additive cluster, and extended tree models) that were not available when Jones and Young conducted their study and see if these models were preferable in terms of goodness-of-fit and interpretability.

Method

Judges and Stimulus Persons

The setting for this study was an urban academic medical center, which comprises three major organizational units: 1) a set of educational programs; 2) a large health service provider; and 3) a general administrative staff, which provided support services for the whole complex. Each

is headed by a Vice-CEO, who reports to a CEO of the center. In addition, there are two staff functions (strategic planning and public relations) which are led by individuals holding the title of Vice-CEO. The 25 persons used as stimuli in the similarity judgments included all of the top two levels of the administrative structure (i.e. the CEO and the five Vice-CEO's), at least a few of the most senior administrators who report to each of the Vice-CEO's, and the senior staff positions (e.g. Assistant to the CEO). The stimulus individuals were selected so as to provide approximately equal numbers of individuals from each of the five major administrative "territories" in the organization. The number of stimulus individuals was limited to 25 in order to keep the number of similarity judgments at a manageable level. (A conventional organizational chart showing the positions of the 25 stimulus individuals is given in Figure 2, and some background information is given as Table 1).

The raters were recruited from among the group of stimulus individuals, and a major criterion was the willingness of the person to perform the tedious task of making the judgments. As a result, the judges do not necessarily represent a cross-section of the stimulus

Table 1:
Stimulus Persons Chart: The Academic Health Center

Plot Symbol	Title	Sex	Race	Comments
1	Administrator of Educational Program III	M	W	A physician and basic science researcher. Has been at the institution for a long time.
2	Assistant Vice CEO for Grants and Conracts	M	W	Recruited relatively recently by the Vice CEO for General Administration.
3	Assistant Vice CEO of Service Programs (Operations)	M	W	A protege of the Vice CEO for Service Programs.
4	Director of Public Relations Function I	F	W	
5	Planning Associate	F	W	
6	Affirmative Action Officer	F	W	One of the "outcasts"- someone who is somewhat on the fringes of the organization. Reports to the CEO, but not close to anyone.
7	Director of Management Systems	M	W	Used to work for the Vice CEO for General Administration, brought by him to the Center.
8	Assistant to the CEO (Appointments)	F	W	A major intermediary for the CEO with other staff, the person who sees him most often. More a secretary than a professional.
9	Administrator of Educational Program I	F	H	Another ""outcast"- not well regarded by others, no close allies.

A	Assistant Vice CEO of Service Programs (Finance)	M	W	Less close to the Vice CEO for Service Programs than the Assistant Vice CEO for Operations, but they are a close-knit trio.
B	Assistant to the CEO	M	W	Senior professional Assistant to the CEO. Came to the institution with the CEO, had worked for him before. Responsibilities for executive recruitment.
C	Controller	M	W	Orthodox Jew
D	Vice CEO for Educational Programs	M	W	A physician. Principal academic officer for the institution.
E	Director of Public Relations Function II	F	B	Another ""outcast"- somewhat embittered, feels unappreciated.
F	Vice CEO of Service Programs	F	W	
G	Planning Associate	M	W	
H	Vice CEO for Public Relations	M	W	Senior public relations officer. Really a staff, rather than operational position. Not highly regarded, in some difficulty at time of study.
I	Administrator of Educational Program II	F	B	Came up through the ranks- student, faculty member, Administrator. Well liked. Made Administrator by the current CEO.
J	Vice CEO for General Administration	M	W	Chief "nuts-and-bolts" administrator. New to medical education and health. Background in public administration.
K	Assistant to the CEO	F	W	Another ""outcast"- never quite been accepted. Rubs people the wrong way.

L	Vice CEO for Planning	M	W	Also an "outcast." Also a staff, rather than operational Vice CEO. Lots of conflict.
M	Assistant Vice CEO for Human Resources	M	H	Relatively recent appointment by the Vice CEO for General Administration
N	CEO	M	W	Chief administrative officer for the institution. A physician.
O	Assistant Vice CEO for General Services	F	W	Worked with the Vice CEO for General Administration before he came to the institution, subsequently recruited by him.
P	Dean of Students	M	B	Been in the institution a long time. A physician. Has more of a faculty role than an administrative one.

Figure 2:
A Partial Organizational Chart for the Academic Health Center

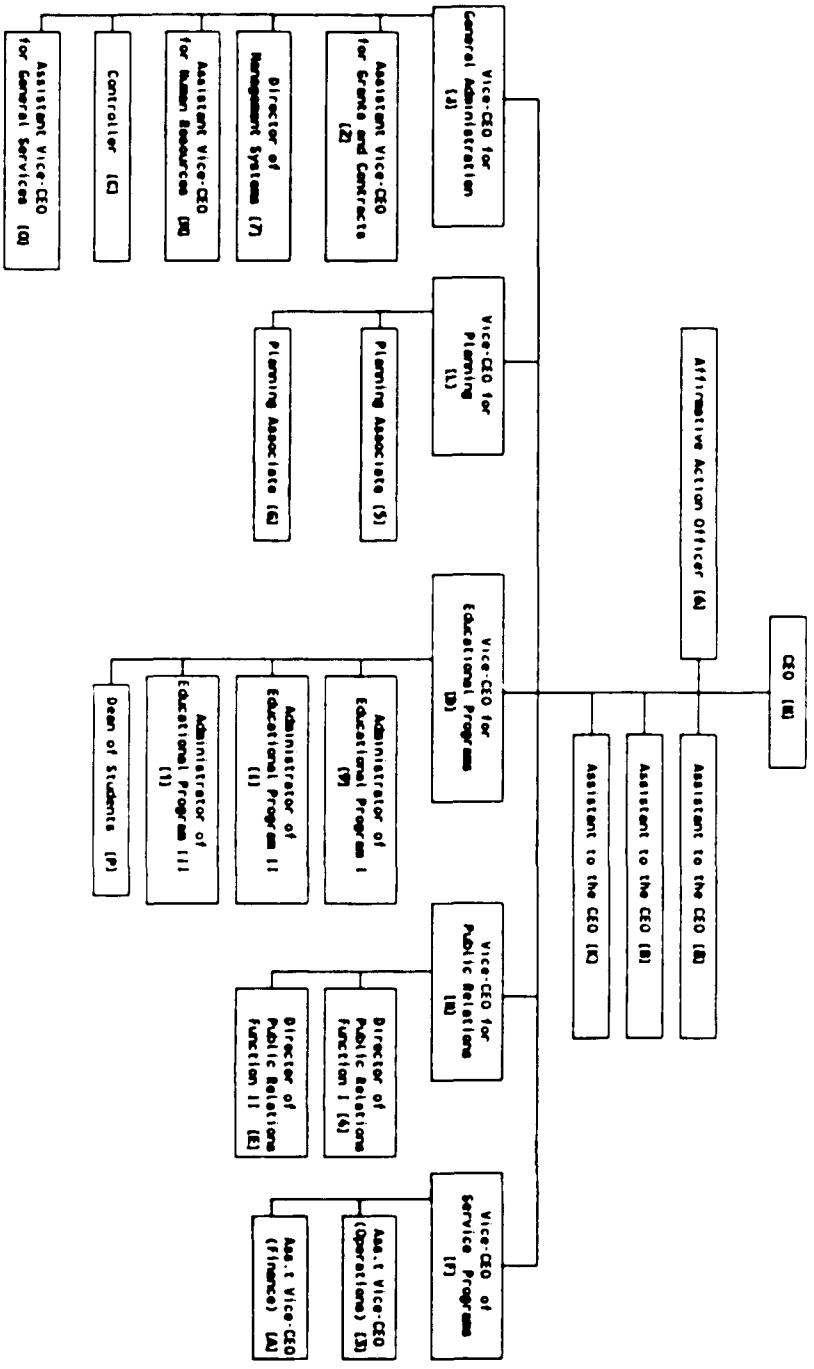

individuals[9]. The judges are indicated in Table 1.

Materials

The ten judges were given a questionnaire consisting of two parts. Part I asked them to indicate "the number on a nine-point scale (from 1= don't know at all to 9= know very well) that best expresses how well you feel you know the person in question", for each of the 25 stimulus persons. Part II presented all of the possible pairings of the stimulus individuals in a "Ross ordering" (Ross, 1934), with the following instructions:

> "In this section, you are asked to rate the similarity of pairs of Medical Center staff, taking into account whatever characteristics you think are relevant. Please circle the number on the nine-point scale (from 1=extremely dissimilar to 9= extremely similar) which best represents the

[9] Interestingly, two of the individuals who were asked to participate in the study refused. Though there is no way to know for certain why these individuals declined, it is noteworthy that both of them were individuals who, in the opinion of the investigator, were experiencing considerable difficulty in taking up their roles in the organization. They are both, in fact, members of the group that is characterized as "outcasts" below. It seems possible, then, that they may have been experiencing a good deal of anger and anxiety, and may have been concerned that their feelings about themselves, their colleagues, and the organization would be revealed through the instrument.

degree of similarity between the two."

The judges rated the similarity between all possible pairs of 25 stimuli. No unidimensional rating scale data were gathered.

Results

The resulting data were then fitted to five statistical models:

1. a multidimensional scaling model, using a individual differences multidimensional scaling program (ALSCAL) similar to INDSCAL;

2. a hierarchical clustering model (TAXON);

3. an additive clustering model (MAPCLUS);

4. an additive tree model (ADDTREE); and

5. a variant of the additive tree model that allows for superimposed cross-cutting categories (EXTREE).

Multidimensional Scaling (ALSCAL)

ALSCAL solutions were obtained in 2 (stress= 0.297, RSQ= 0.570), 3 (stress= 0.212, RSQ= 0.628), 4 (stress= 0.169, RSQ= 0.647), and 5 (stress= 0.145, RSQ= .663) dimensions[10]. Graph 1 presents the total variance accounted for each of these MDS solutions. Beyond three dimensions,

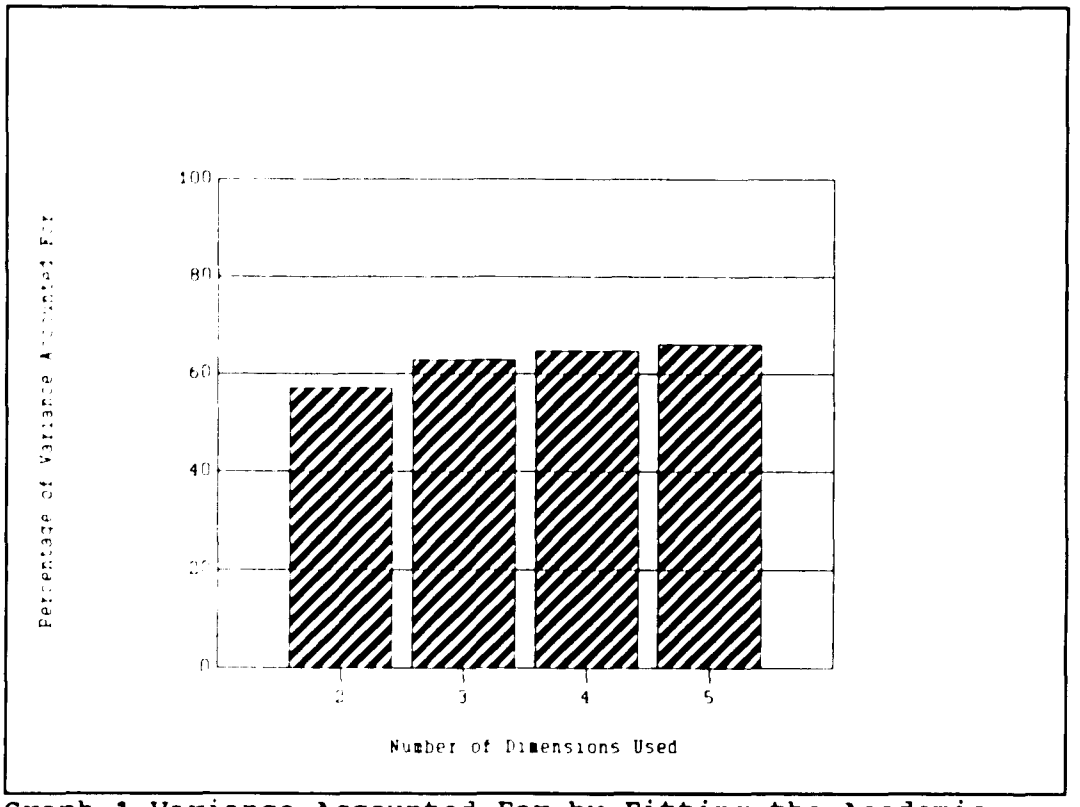

Graph 1 Variance Accounted For by Fitting the Academic Health Center Data to the MDS Model

[10] Stress is a measure of goodness of it- the lower the stress, the better the fit. RSQ is the percentage of similarity variance accounted for by the MDS output configuration.

the gains in variance accounted for are small, suggesting that the three dimensional solution probably contains most of the information that the spatial model is capable of extracting from this data.

All four solutions were generally consistent. The common structure is most easily visualized in the three-dimensional solution, which is presented in Figures 3 and 4.

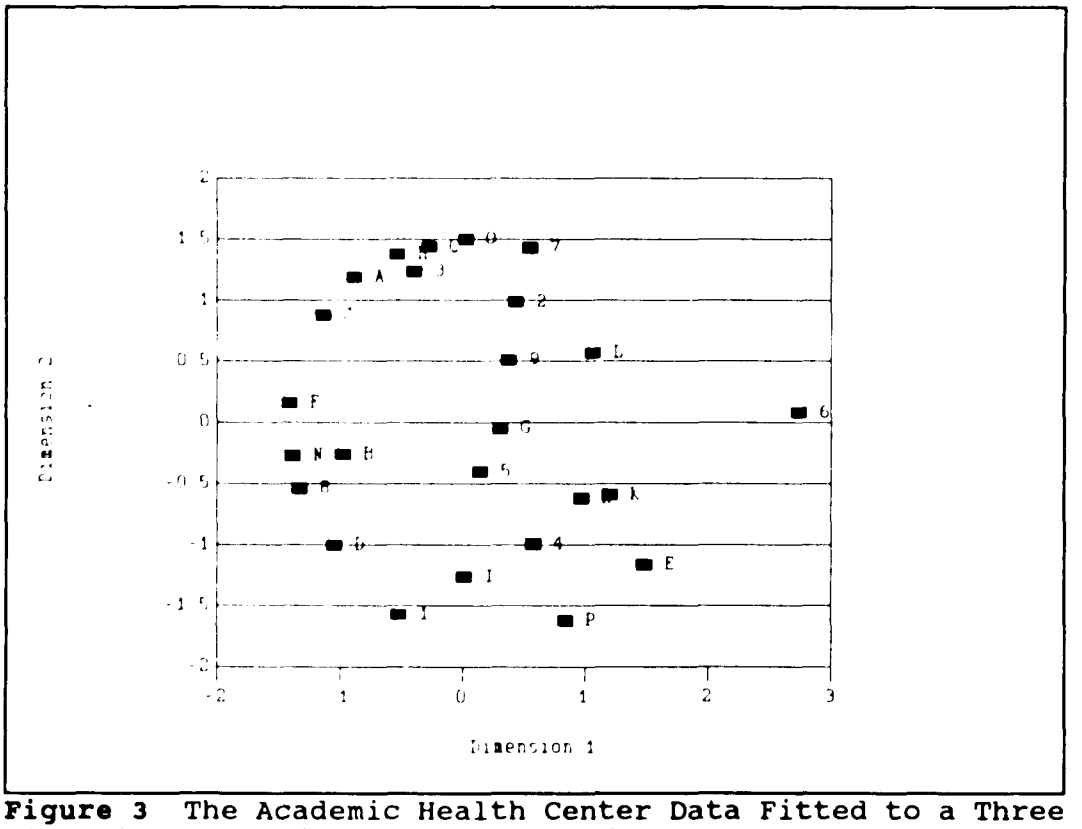

Figure 3 The Academic Health Center Data Fitted to a Three Dimensional Spatial Model (Dimensions 1 and 2)

Dimension 1 appears to correspond roughly to status. In a general way, this dimension resembles the vertical dimension of the standard organizational chart. There are several individuals, however, whose titles (including the Vice CEO's for Planning - Plot Symbol L and Public Relations- Plot Symbol H) suggest a high status but who are placed relatively low on Dimension 1.

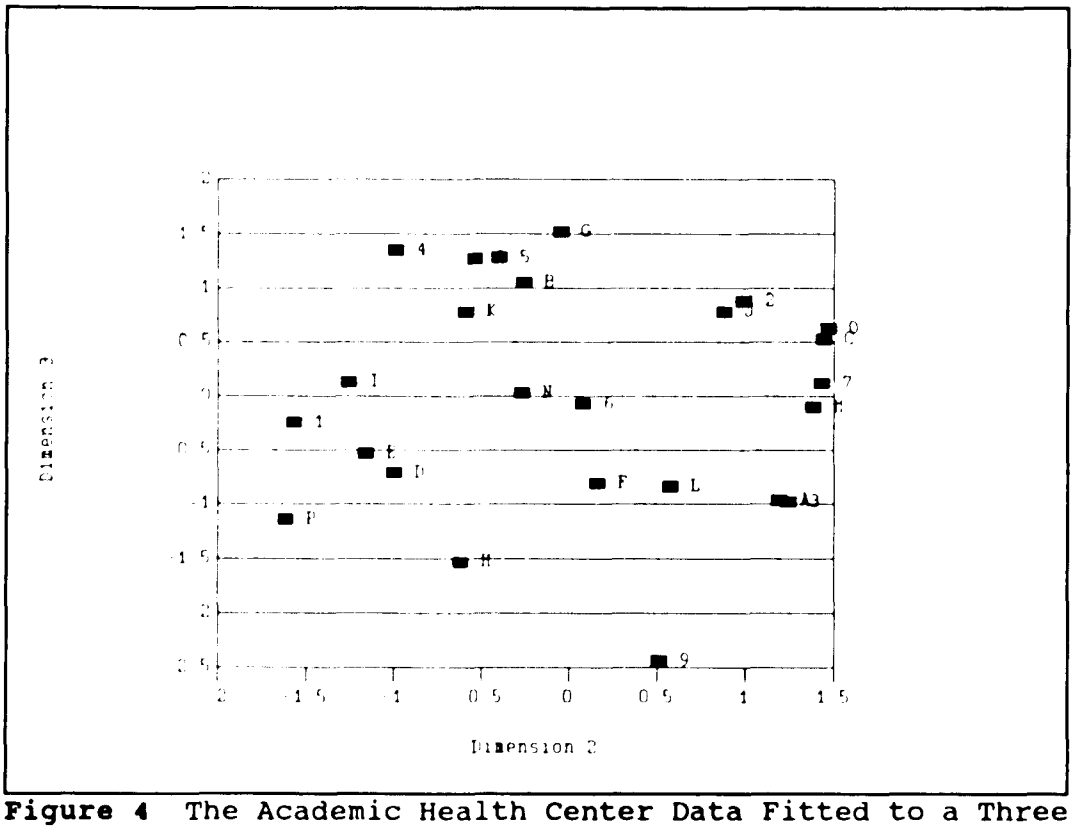

Figure 4 The Academic Health Center Data Fitted to a Three Dimensional Spatial Model (Dimensions 2 and 3)

In the absence of confirmatory rating scales, it is impossible to come to any firm conclusions, but two considerations may be relevant. First, neither of these two Vice-CEOs were in charge of major operational units. Instead, they were responsible for what might be better described as staff functions. It may be that the raters did not see them as being as powerful as the other Vice-CEOs, or, perhaps, even as powerful as some of the Assistant Vice-CEOs (most of whom have far more staff reporting to them). Second, both of the Vice-CEOs and the Assistant to the CEO who were placed low on Dimension 1 were having some difficulty in the organization at the time the data was collected, and may well have been seen as being held at some distance by the CEO and the other Vice-CEOs.

Considered independently, interpretations for the other two "dimensions" were not immediately obvious. However, when the two dimensions are used to define a two-dimensional space (see Figure 4), a plausible interpretation readily emerges. Four "clusters", roughly corresponding to the four quadrants of the plane, can be seen, which appear to coincide with the three major operational areas of the institution (academic affairs, general administration, and hospital administration) and a fourth category that might be thought of as "staff." The CEO's location, almost at the dead center of the coordinate axes, equidistant from all of

the clusters, seems appropriate.

Thus, we can visualize a three-dimensional arrangement as follows (Figure 5): a vertical dimension, corresponding to status, with the CEO at the top, and four clusters, each differentiated according to status, three of them "hanging" from their Vice-CEOs. This is a very creditable representation of the standard organizational chart.

Following this line of thought, there are grounds for a non-spatial interpretation of this data. It could be argued that the computer, in attempting to fit "cluster data" to a spatial model, has produced a pyramidal configuration (in three-dimensional space) that is an attempt to represent the following "cluster" structure: the four corners of the base of the pyramid are the four functional clusters described above; the top of the pyramid is another group (call it "the top management group") consisting of the CEO, the three operational Vice-CEOs, and two of the Assistants to the CEO. The three Vice-CEOs (and perhaps the two Assistants) are members of both the top management group and of their respective operational unit clusters, and their spatial locations in the MDS solution are a way of "expressing" this

Figure 5:
A Three-Dimensional Representation of the MDS Configuration Obtained by Fitting the Academic Health Center Data to a Spatial Model (ALSCAL Algorithm)

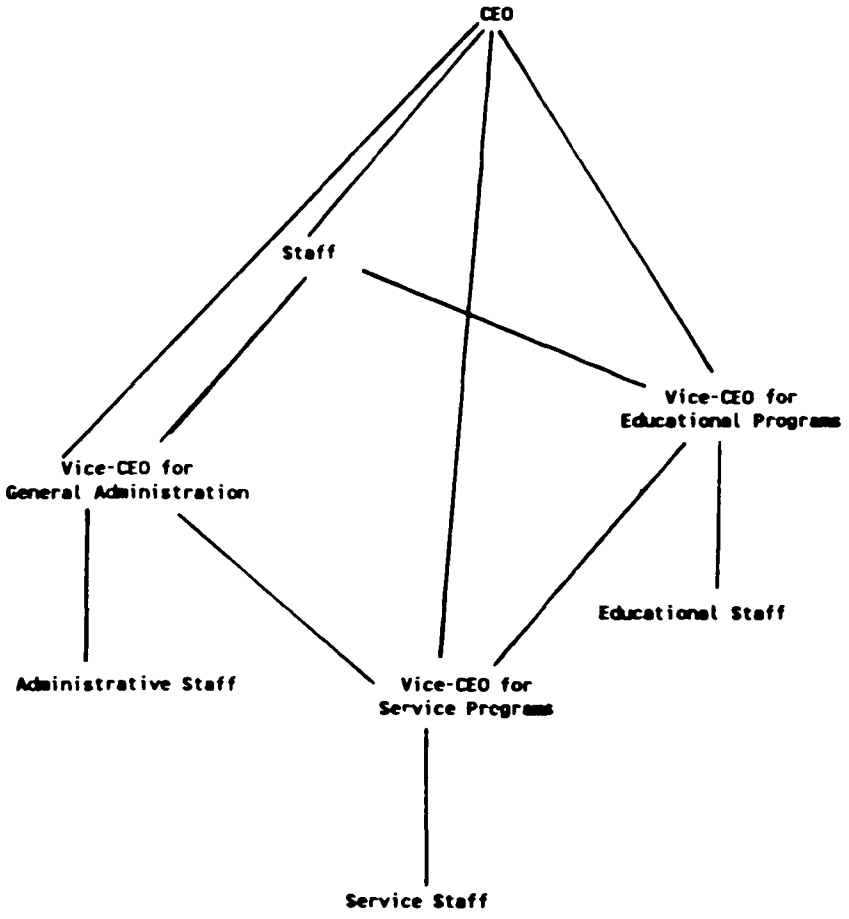

dual membership[11].

Thus, this analysis produces a fairly good reconstruction of the organizational chart, but does not allow us to go much beyond that. There is no evidence that the multidimensional scaling technique is sensitive enough to allow us to detect the intrusion of task-irrelevant sentience groupings. It is possible, however, that the use of a statistical analysis that is more specifically appropriate to our model might produce a better result[12].

Hierarchical Clustering (TAXON)

In order to pursue the idea that some of the variance in the similarity judgments might be contained in a clustering structure which could not be brought out by a MDS analysis (i.e. that some of the "clustering information" was discarded as the data was fit to a MDS model), a

---

[11] This structure also appears, sometimes in collapsed form, in the four-dimensional and two-dimensional arrangements, suggesting that it is not adventitious.

[12] It may be that, in fitting the data to a spatial model, the computer was required to construct a two-dimensional spatial analogue of what we suspect may be an inherently "cluster-type" structure. The "clusters" that are created are "smudged"- that is, individuals that are members of the same cluster, and which should therefore have exactly the same value, are spread out in the plane. It would seem that this "smudging" must "soak up" some variance that might well carry some of the other aspects of sentience.

hierarchical clustering analysis was performed on the same data set, using a program called "TAXON."

As discussed above, the hierarchical clustering model is based on the assumption that stimuli occupy locations in nested categories, and that the psychological distance between two stimuli is determined by how high in the category structure one needs to go to find a category that encompasses both stimuli[13].

The result of the TAXON analysis are given in Figure 6. The overall correlation between the hierarchical structure given there and the original similarity data is .70, indicating that this structure is capable of accounting for about half of the original variance, not as good a result as that achieved by MDS. Nonetheless, the results are readily interpretable, and display aspects of the organization that were obscured or discarded in the MDS output.

The TAXON structure should be interpreted as follows. Starting from the lower left corner of the chart, the structure splits the group into two large clusters. One

[13] An analogy might be mailing addresses- house, street, zip code, town, state, and country form a structure of nested categories, and distance could be thought of as a function of the lowest level category into which two addresses both fall.

Figure 6:
The Academic Health Center Data Fitted to a Hierarchical Clustering Model (TAXON Algorithm)

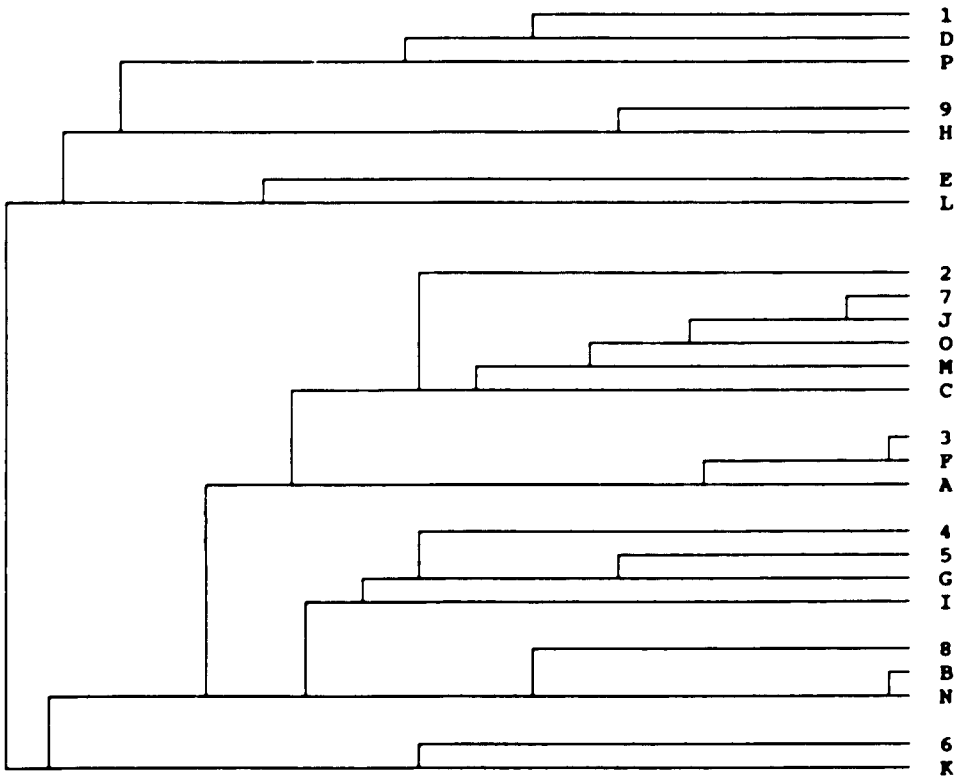

(containing stimulus persons 1, D, P, 9, H, E, and L) corresponds roughly to the educational services cluster, while the other contains the administrators and staff. In both branches, the next bifurcations separate the individuals whose locations seem to be somewhat problematic and who were placed in anomalous locations by the MDS analysis (9, H, E, and L on one branch; 6 and K on the other). Thus, the interpretation that these individuals are people "that the organization doesn't know what to do with" or "who don't fit in" is supported. Alternatively, they may be individuals who are "marked" by salient features that lead them to be categorized in ways that are different from the categorization strategy used for the rest of the set. In the discussion to follow, these individuals will be referred to as the "outcasts."

In the lower set of branches, the next bifurcation splits the staff administrators centered around the CEO (4, 5, G, I, 8, B, and N) from the line managers in both the Service Program area and the general administration. The first of these clusters is then partitioned further into the CEO's Office staff and others, and then the CEO's Office staff is further divided into the CEO and the senior assistant (who came to the institution with the CEO) on the one hand, and the CEO's executive secretary on the other. The line administrative cluster is partitioned in a similar

manner. First, the service program and general administration groups are distinguished, and then the fine detail of the partitions that follow also mirror the length of association (and prior associations) of the staff. It might be said that the clustering output is well suited to the representation of an executive's "inner circle." In the MDS analysis, these details did not emerge, and this may be the strength of the hierarchical clustering analysis- that it represents the "fine structure" of the clusters in the organization better than MDS can.

On the other hand, several important aspects of the "real organizational structure" are lost in the hierarchical clustering analysis:

1. status is invisible;

2. the group of senior managers (The CEO and his Vice CEOs) that the MDS was able to represent has been broken up, probably resulting in an "overstatement" of the distances between the senior staff[14].

---

[14] This is a consequence of the hierarchical clustering requirement that there be no overlapping clusters- which does not allow the Vice-CEO for General Administration, for example, to be simultaneously a member of the "executive group" and a member of a cluster that comprises his staff. As a result, the program is forced to choose to represent one, and discard the other. The gain in

3. the centrality of the CEO, and his "equidistance" from each of the clusters, is not evident.

On the basis of these results, several conclusions seem reasonable. First, it appears that MDS is better at representing the "real" cluster structure than hierarchical clustering is at representing variables such as status that MDS represents well. This seems reasonable in light of the fact that MDS accounts for somewhat more of the variance than the hierarchical clustering analysis. MDS, however, does not seem to be particularly good at representing the fine structure of the groups, which hierarchical clustering may do well. Second, both approaches seem to be capable of identifying the "outcasts"- though it is not clear how a reliable "index of outcast-ness" could be developed. Third, though both approaches are capable of revealing aspects of the "real structure", neither is capable of revealing it all. Thus, the idea that the "real structure" is a hybrid seems to be generally supported. Overall, it would seem that the validity of the general approach- using similarity data as the basis for reconstructing the "map in the head"- is supported.

"variance-accounted-for" is greater if this Vice-CEO is placed in the same cluster as his (relatively numerous) staff members.

Additive Clustering (MAPCLUS)

The data was fitted to additive clustering models using one, two, three, four, five, six, seven, eight, nine, ten, eleven, twelve, fourteen, and sixteen clusters. In many cases, two or three trials were performed. Graph 2 gives the variance accounted for statistic for each of the solutions.

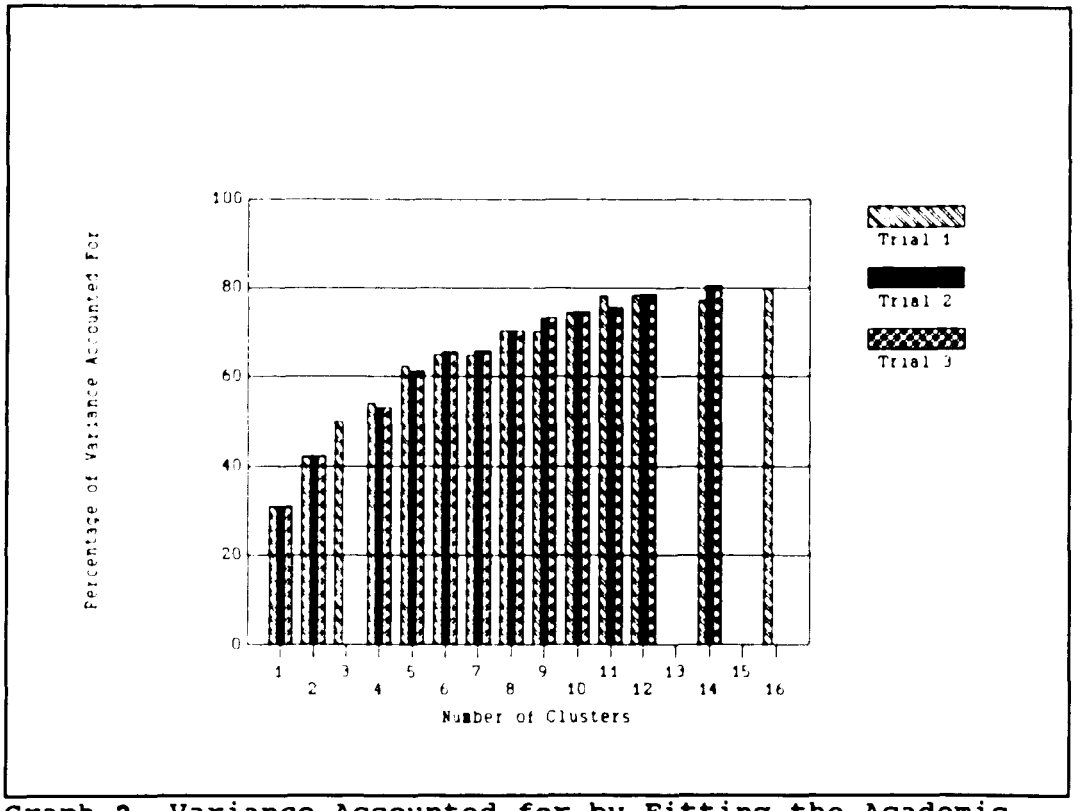

Graph 2 Variance Accounted for by Fitting the Academic Health Center Data to the Additive Cluster Model (MAPCLUS Algorithm)

The discussion here will focus on the solutions using eight clusters or less. Beyond eight clusters, the increments in variance accounted for are relatively modest and, more importantly, there are no discernible improvements in the interpretability of the clusterings that are produced. Furthermore, beyond twelve clusters, the MAPCLUS algorithm frequently failed to arrive at a solution, suggesting that too many clusters were being used.

Three clusters emerged consistently throughout the MAPCLUS analyses:

1) a cluster containing the three Service Program administrators, which was the first to emerge (as the number of clusters was increased) and which was consistently marked by a high weight, indicating a high degree of salience;

2) a cluster (also with a relatively high weight) containing the all of the general administrators, except the Assistant Vice CEO for Grants and Contracts; and

3) a cluster containing the CEO, the Vice-CEO for Service Programs, the Vice CEO for General Administration, and the two "non-outcast"

Assistants to the CEO.

The appearance of the first two clusters confirms the results of the other methods. The third cluster is important in two respects. First, from a methodological point of view, it appears to correspond to the "senior management group" whose existence had been hypothesized, but which the hierarchical clustering approach had obscured. This is not surprising in that the additive clustering model is well suited to bring out a cross-cutting cluster such as this, but other methods (see the discussion of the EXTREE model below) can produce the same results much more easily.

Second, it is noteworthy that when the "senior management group" emerges, it does not include the Vice CEOs for Planning, Public Relations, or Educational Programs, and does include two of the CEO's staff. The exclusion of the two "outcast" Vice CEOs (who were, as well, placed low on the "status" dimension of the MDS solution) is not very surprising, but the failure to include the Vice-CEO for Educational Programs, who was not obviously "outcast" is more interesting. It suggests that the Vice-CEO for Educational Programs, though holding a powerful position, is, perhaps, not perceived to be a member of the CEO's inner circle in the way that the other members of this cluster are.

Overall, however, the MAPCLUS results are disappointing. The method is computationally cumbersome, lacks the convenient graphical representation that would make it easy to interpret, and does not produce especially valuable results. It may, however, be useful in confirming the results of other methods.

Additive Tree Models (ADDTREE and EXTREE)

Three variations of this general class of models will be discussed here:

1. an additive tree model (ADDTREE);

2. an extended tree model, which allows for the superposition of cross-cutting clusters on top of an additive tree structure (EXTREE); and

3. a variation in which the "organizational chart" structure is specified as an initial configuration for an EXTREE analysis.

ADDTREE

The result of fitting this data to the additive tree

model is shown as Figure 7. This configuration accounted for 68.14% of the total variance. Several features are worth noting.

For one thing, the cluster structure in the ADDTREE analysis bears a strong resemblance to the hierarchical clustering analysis described above. The CEO's Office, Service Program Administration, and General Administration clusters all emerge strongly, and appear to be organized around their leaders, just as they did in the TAXON analysis. The Educational cluster is entangled with some of the Public Relations staff, but the core of physicians around the Vice CEO for Educational Programs is also distinctly visible. Many of the individuals whose location in the TAXON and MDS analyses did not appear to correspond to their formal title in the organization (the "outcasts" discussed above) also appear to be "out of position" in the ADDTREE analysis. For example one of the Assistants to the CEO ("K") is located far outside the CEO's Office cluster. Thus, the results of the two approaches are consistent in suggesting that some of the organizational clusters (e.g. the General Administration cluster) are more clearly defined and bounded in the minds of the members of the group, while others (e.g. Public Relations) appear to "agree" be far less cohesive. One way to interpret this is to suppose think of the regularly recurring clusters as consisting of

Figure 7:
The Academic Health Center Data Fitted to an Additive Tree Model (ADDTREE Algorithm)

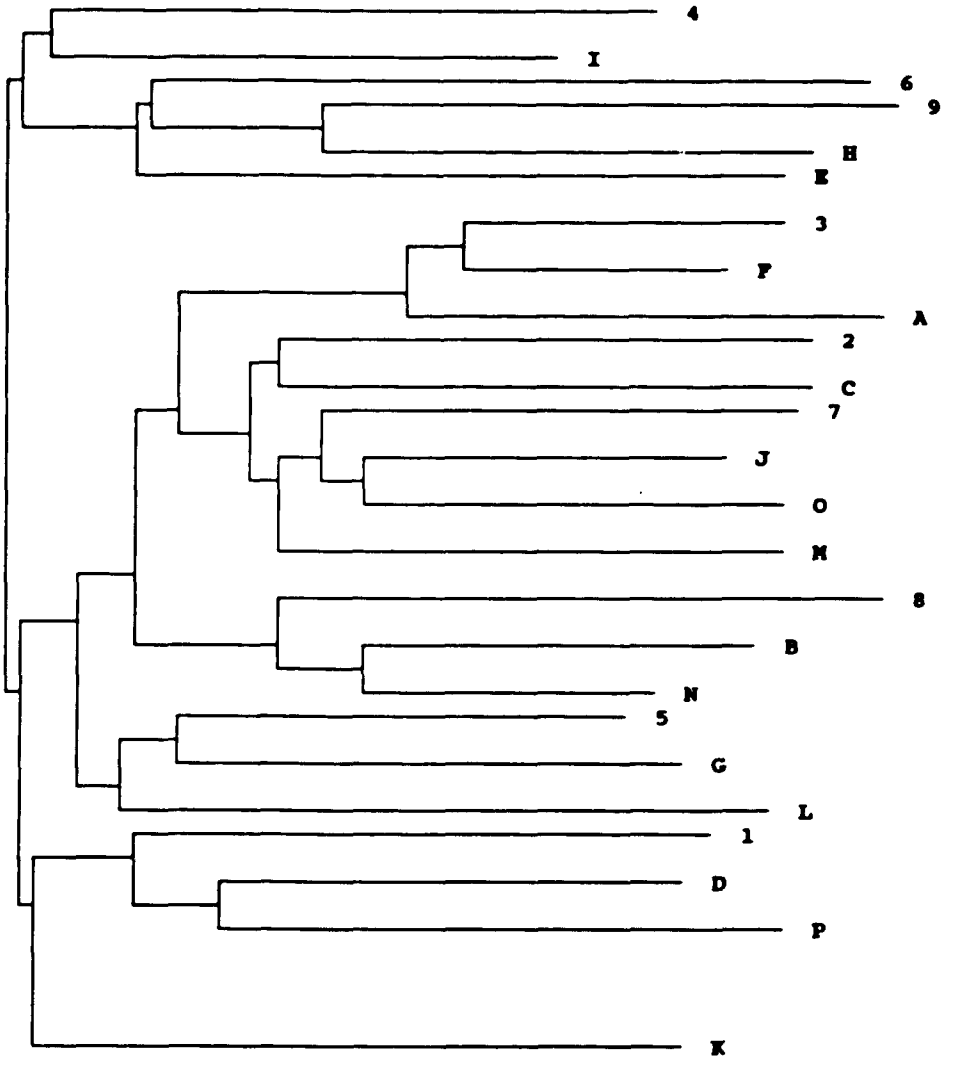

individuals who are strongly knit together by "similarity bonds", while those clusters that are more transient are, we might say, more "friable"- that is, the "similarity bonds" are weaker. This consistency would appear to indicate that these facts are inherent in the similarity data and not artifacts of the statistical procedures used.

Furthermore, the ADDTREE model, unlike TAXON, allows for variations in the length of the branches connecting the stimulus individuals to the root of the tree, and these variations can be interpreted as indications of the degree of "distinctiveness" of that individual. That is, the longer the branch, the greater the difference between that individual and the rest of the group and the shorter the branch, the stronger that person's connections to the rest of the group. This feature of the model is useful in two respects.

First, one can observe that many of the individuals with the longest branches (the Affirmative Action Officer, for example) are among the group of "outcasts." Graph 3 gives the ADDTREE extensions (in z-scores) shows that the individuals with exceptionally long extensions (the Affirmative Action Officer, one of the Assistants to the CEO, and the Planning Vice-CEO) tend to be among the outcasts. (Similar results could be obtained for the EXTREE

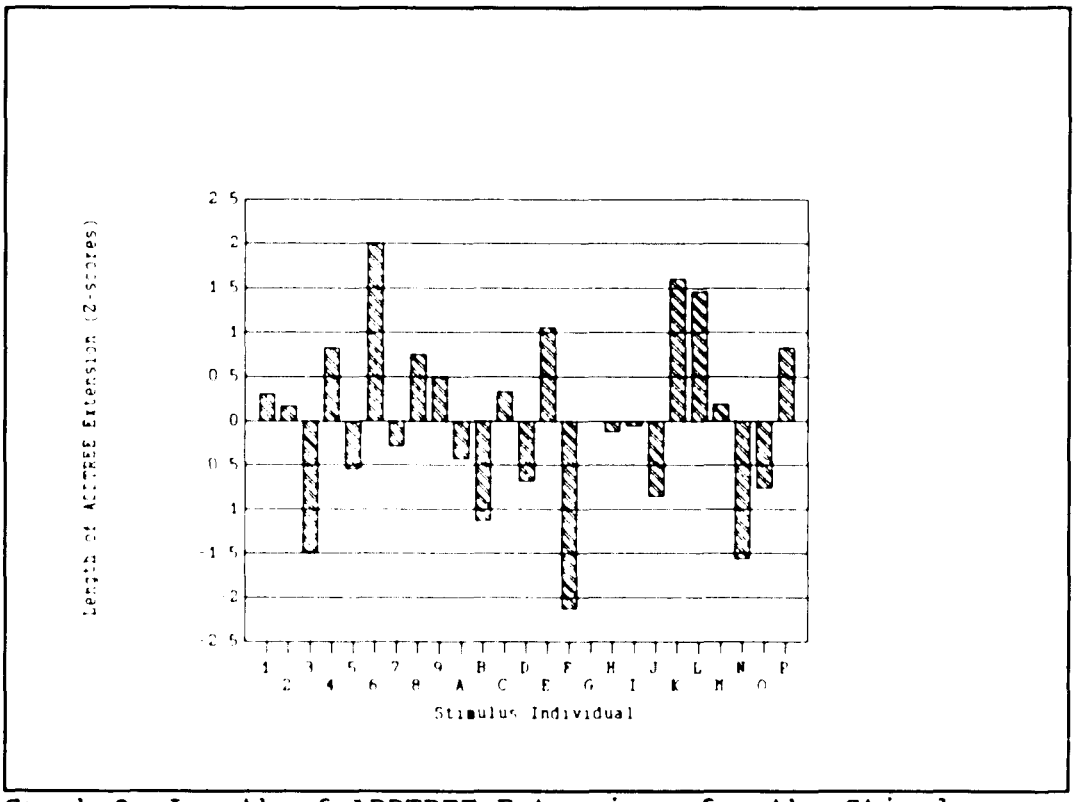

Graph 3 Length of ADDTREE Extensions for the Stimulus Individuals: The Academic Health Center Data

extensions, see below.)

Second, one may observe that within the well-defined clusters, the leader is also the individual with the shortest branch. In the Service Program cluster, for example, the Vice-CEO's branch is distinctly shorter than either of the other two members of the cluster[15]. It would

[15] Even the exceptions seem to prove this rule, as we note that the Planning Vice-CEO, one of the "outcasts", has the longest branch in his cluster. The Planning cluster is atypical in another respect, in that its leader is not at the core of the set of subclusters, but in an outer layer.

seem reasonable to interpret this as an expression of the Vice-CEO's broader responsibilities and of her affiliation with the overall institutional management. Thus, it would appear that the length of branch expresses somewhat the same sort of information that the "status" dimension of the MDS analysis did. Furthermore, it would seem reasonable to suppose that this additional parameter may explain a good deal of ADDTREE's ability to account for a substantially greater percentage of the variance than TAXON could. All in all, there seems to be no reason to prefer TAXON to ADDTREE, and good reason to prefer ADDTREE.

EXTREE

The output for the EXTREE analysis is shown in Figure 8. This combination of a tree structure and a set of twelve superimposed clusters accounted for 76.54% of the total variance. The overall tree structure is generally similar to the ADDTREE structures.

Two observations can be made about the effect of adding the cross-cutting clusters to the model. First, the additional parameters produce a very modest increase in variance accounted for. In general, as will be discussed below, once one of these models has reached the point of accounting for about 65-70% of the total variance,

Figure 8:
The Academic Health Center Data
Fitted to an Extended Tree Model (EXTREE Algorithm)

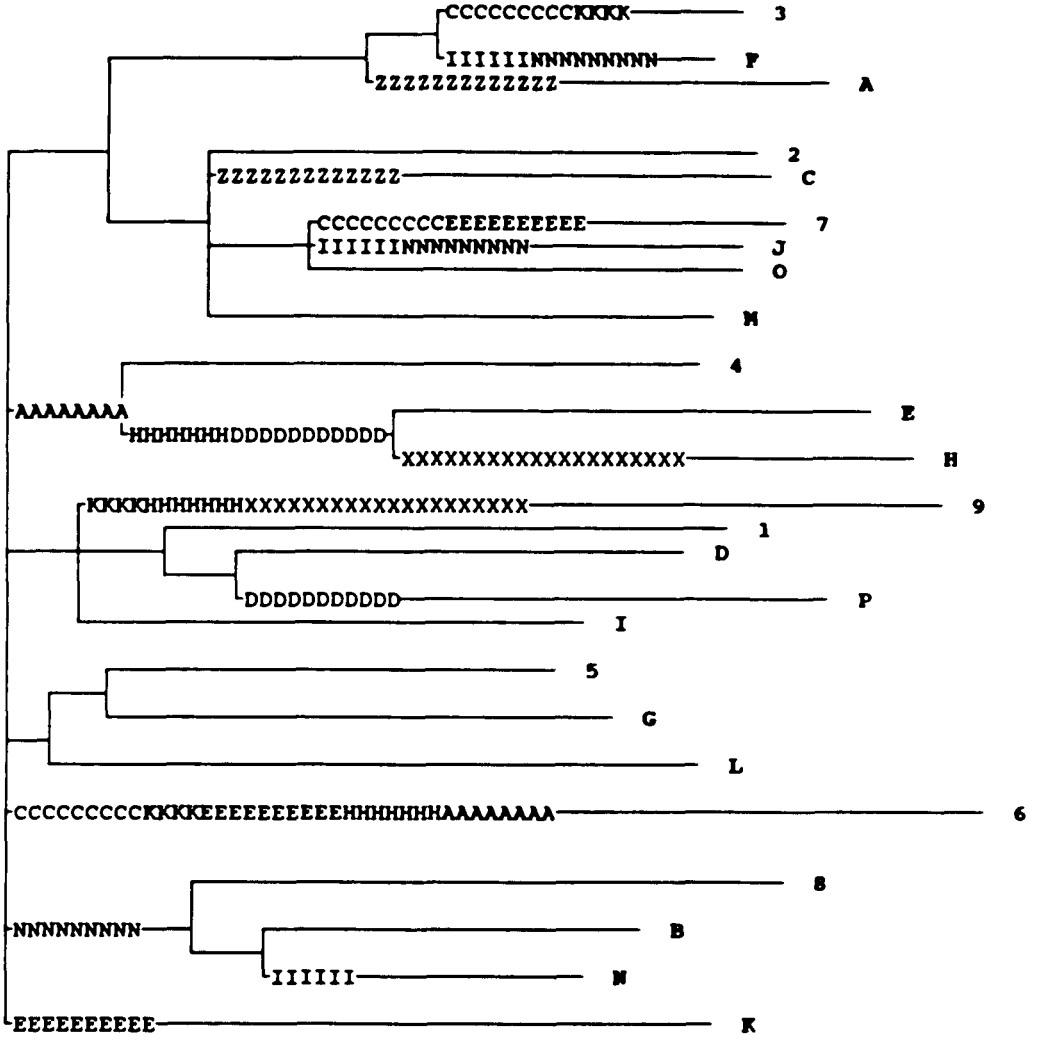

additional parameters produce very little improvement in accounting for the variance in this set of data.

Second, only a few of the superimposed clusters appear to have obvious interpretations. In fact, the "outcast" individuals are very well represented in the additional clusters, which suggests that much of the (non-error) variance left unaccounted for (after the basic tree structure has been taken into account) is associated with the outcasts. Furthermore, since the "outcasts" do not represent a cohesive group, but rather individuals who have been "pushed out" of the locations that they would be expected to occupy, no sensible groupings can be formed by collecting them together.

Some of the superimposed clusters, however, do seem to have plausible interpretations. Clusters I and N seem to correspond to the senior management group that the hierarchical clustering model tends to break up. As noted above, ADDTREE, unlike TAXON, allows this group to reemerge through the length of the branches, but it seems possible that there is still some residual variance attached to the senior management group that can be expressed through the superimposed clusters.

Cluster Z, composed of the Controller and the Associate

Director for Finance, can be interpreted as representing a distinctive area of interest that cuts across the functional divisions within the institution. It appears consistently throughout the EXTREE analyses, suggesting that it is not a statistical artifact.

EXTREE with Pre-Specified Clusters

In this procedure, the similarity data is first fitted to an additive tree model, using ADDTREE, and cross-cutting clusters are then superimposed onto the model, using the EXTREE procedure. Unlike conventional EXTREE, however, the investigator can specify the membership of the clusters that are to be added, and can add them one at a time and then measure the improvement in goodness-of-fit due to the new cluster. What this amounts to, in a sense, is a test of the hypothesis that there is still, after the data has been fitted to a tree structure, some residual component of the overall similarity variance that can be explained by the cross-cutting cluster. For example, if one felt that gender was a "hot" category in the organization at the time that the data was gathered, one might expect that: 1) gender would exert a detectable influence on the similarity judgments; 2) therefore, a detectable portion of the similarity variance would be unexplained by the tree structure; and the variance accounted for would increase if

a cross-cutting gender cluster was added to the tree model. In this way, the procedure can be used as a means of detecting the strength of the influence of task-irrelevant sentience groups on the commitment of the members to the task structure.

In this study, six cross-cutting clusters were tested in this manner:

1. the "Senior Management Group": The CEO, Vice-CEO for Service Programs, Vice CEO for Administration, and two Assistants to the CEO;

2. a group consisting of five of the "outcasts": one Assistant to the CEO, the Vice CEOs for Planning and Public Relations, the Affirmative Action Officer, and the Administrator of Educational Program I;

3. the females;

4. the minorities (Black and Hispanics);

5. the minorities plus the Affirmative Action Officer; and

6. the Controller and the Associate Vice-CEO for Finance.

The ADDTREE configuration shown in Figure 7 (see page 92) was used as the starting point for this analysis. Table 2 presents information pertaining to the impact of the added clusters on goodness-of-fit. None of the clusters appear to contribute very much, with the exception of the management group, suggesting that it may have a very real presence in the minds of the raters, especially since, as noted above, some of the commonality of this group is represented by various features of the ADDTREE configuration. The cluster containing the two financial individuals is the next strongest. The task-irrelevant clusters (e.g. the women, the minority groups) do not display visible influences in this data, and we must therefore conclude that the data offers no support for the hypothesis that these subgroups are prominent in organizational life, at least in comparison to the task-relevant groups.

Table 2:
The Effect of Adding Cross-Cutting Clusters
to an ADDTREE Model:
The Academic Health Center Data

Cluster Added to Model	Increase in Variance Accounted For	Decrease in Stress	
		Formula #1	Formula #2
The Senior Management Group	2.79%	.0030	.0307
The "Outcasts"	0.26%	.0001	.0014
The Women	0.30%	.0001	.0011
The Minorities	0.51%	.0001	.0027
The Minorities plus the Affirmative Action Officer	0.48%	.0002	.0032
The Controller and the Vice-CEO for Service Programs (Finance)	0.82%	.0010	.0104

Discussion

At this point, it may be useful to consider the two issues raised at the outset of the pilot study.

First, do the results of the pilot study tend to provide support for the idea that similarity ratings can be used to reconstruct mental representations of social space, as suggested by the Jones and Young study? It would seem that they do. Every one of the techniques that were applied to the data from the pilot study produced configurations that were recognizably similar to the basic structure of the organization (as described by the organizational chart). Furthermore, even the discrepancies between the chart and the output configurations appeared to be interpretable, the best example being the "outcasts"- a set of people who could be said to be having difficulty taking up their roles in the organization, who could be identified by virtue of their anomalous location in the output configurations.

The results of the pilot study make it possible to advance and support five hypotheses about this organization:

Hypothesis #1 (H1): An individual's similarity relationships, and therefore his/her location in "organizational space", are primarily defined by status and functional role.

This is a version of the conclusion offered by Jones and Young. It is generally supported by all of the analytic techniques discussed above:

Multidimensional Scaling (ALSCAL)- the algorithm produced a three-dimensional spatial arrangement in which one dimension corresponded roughly to status in the hierarchy and the other two established four clusters corresponding to functional area;

Hierarchical Clustering (TAXON)- provided confirmation of the functional category structure, though the model precludes the emergence of the status aspect;

Additive Clustering (MAPCLUS)- supported the existence of a senior management group which roughly corresponded to the higher stratum of the MDS status dimension as well as the Administration and Service Program Clusters;

Additive Tree Models (ADDTREE and EXTREE)- produced a functional category structure very similar to those produced by TAXON and also, through the fine structure of the categories and the length of the branches, produced indications that corresponded to the status dimension.

While the data and analysis from the pilot study is generally consistent with the results from the Jones and Young study, there are several differences. First, where Jones and Young treat the functional aspect as a continuous spatial dimension, the discussion above suggests that it may be much more useful to treat it as a structure of nested and overlapping categories. Analyses based on this assumption are capable of revealing a "fine structure" within the categories that is beyond the reach of the spatial model. Second, the pilot study suggests some alternative approaches to understanding the status dimension. For example, both the additive cluster and additive tree analyses identified a small group around the CEO which may be thought of as a senior management group, and it may be that status can be regarded as a function of membership in groups that "revolve around" powerful individuals.

Hypothesis #2 (H2): There is a small circle of individuals around the CEO which includes two of the Assistants to the CEO, the Vice CEO for General Administration, and the Vice-CEO of Service Programs, and does not include the Vice CEOs for Public Relations, Planning, or Educational Programs.

H2 is at least partially supported by all but one of the analytic techniques. It is most explicitly supported by the MAPCLUS and EXTREE analyses, which independently constructed this small subgroup from the similarity data. A weaker version of H2, that the Vice CEOs for Public Relations and Planning are not members of the group around the CEO, is supported by both ADDTREE (in which these individuals had long extensions) and ALSCAL (in which they were placed low on the status dimension, at the opposite pole from the CEO). The assumptions that underlie TAXON make it very difficult for support for this hypothesis to emerge from that analysis, and TAXON's failure to confirm does not weigh heavily against the hypothesis.

Hypothesis #3 (H3): Some individuals (certainly one of the Assistants to the CEO, the Vice CEOs for Public Relations and Planning, and perhaps, to a lesser degree, some others) are having difficulty taking up their roles in this organization.

All of the analyses except MAPCLUS identified, in different ways, the "outcasts." In general, they can be thought of as people who are not placed, by the group of raters, in the locations where their roles would lead us to expect that they should be. These analyses do not offer any clue as to whether these difficulties have arisen because the individuals were unable to take up their roles, or not allowed to take them up.

Hypothesis #4 (H4): The boundaries of some of the functional clusters (the Service Program, General Administration, and the CEO's Office) are more sharply defined than those of others (Planning, Public Relations, Educational Programs).

All of the analyses support this hypothesis. The three sharply defined clusters emerge consistently, which the others are either much more blurred or fail to emerge. This suggests that the strength of the identifications of the members with these groups is especially strong. As we have seen above, this may or may not be desirable, depending on the circumstances.

Taken together, H2, H3, and H4 may shed some light on the role of the "outcasts." It may be that one of the tasks of a leader of a work group (for example, a Vice CEO in this

organization) is to articulate a "sense of identity" for the group which can generate commitment from the members. If they succeed, the group will have a distinct psychological boundary and high coherence. If they fail, or are not allowed to succeed, the group will be more "friable." We might also expect that groups which have strong boundaries would display high intragroup cooperation. (They might also have some difficulty cooperating with other groups in times of trouble, if they had a tendency to displace blame for problems onto those other groups.)

Hypothesis #5 (H5): Task-related factors (e.g. status, functional role) predominate over non-task-related factors (e.g. race, gender) in shaping the experience of sentience in this organization at the time that this data was collected.

The evidence for the second part of this hypothesis is indirect. Task-related factors emerged in all of the analyses and non-task related factors did not emerge significantly, even when they were explicitly "encouraged" as when specified clusters were tested through the EXTREE analysis.

Intriguing as these hypotheses are, they cannot be confirmed at this time, as there is no external, objective

information that would serve as a test. This is a weakness of the pilot study design which will be addressed in the design of the study proposed here.

A second purpose of the pilot study was to compare the performance of a wide range of models of similarity and their associated algorithms for the analysis of similarity data. Much discussion of this issue has been included in the detailed review of the results of each algorithm above. Overall, it would appear that for this data, the two additive tree models (ADDTREE and EXTREE) are to be preferred. They provide a convenient, informative output that is easy to read, and account for as good a percentage of the variance as any of the other methods. There is one important qualification, however. The additive tree model is well suited to bringing out the kind of structure that the organizational chart is intended to represent, and it may be that this model is only suitable in organizational situations where the categories that are depicted in the chart can be expected to have a strong influence on the identifications. In different situations (for example, a new organization in which individuals are just taking up their roles, or an organization in such difficulty that role identifications are breaking down) other models may prove to be more capable of capturing the structure that it latent in the similarity ratings.

The major limitation of the pilot study was the lack of objective information (such as the unidimensional rating scales used by Jones and Young) about the functioning of the organization that would allow correlations to be made between the output of the analyses of the similarity data and the behavior of individuals in the organization. As a result, there is little evidence that the interpretations that have been suggested for the discrepancies between the organizational chart and the results of the analyses are correct. This limitation was addressed in the studies described below.

IV. THE SECOND STUDY: AN INTERDISCIPLINARY TEAM

Introduction: The Aim of the Study

The aim of this study was to provide additional support for the validity of the approach described above. This was done in two ways. First, the members of the group that under study had an opportunity to review and confirm or disconfirm the hypotheses that are generated by the analysis of the similarity data. As noted above, this procedure is comparable to Jones and Young's use of unidimensional rating scales to support their interpretations of the dimensions of their three-dimensional MDS solution. If the approach is valid, we would expect that the subjects would confirm the hypotheses generated on the basis of the similarity data.

Second, the method was applied to a specific setting (an interdisciplinary health care service team) which has been analyzed from the point of view of sociotechnical theory. To the degree that the predictions of the theory and the patterns revealed in the data correspond, both the theory and the approach will be supported.

Method

Subjects

The subjects of the study were the eighteen members of an interdisciplinary health care group at an urban academic health center. Not all subjects participated in all stages of the data collection (see below). Table 3 gives background on the members of the group. The focus of the group is a disease syndrome that has relatively recently become a matter of special concern at the institution and the group has been in existence for less than ten years. It is also noteworthy that interdisciplinary groups of this type, that cross the traditional departmental boundaries of academic medicine, are a new phenomenon in this institution, and that there are, therefore, few established models for how they should be organized, how they should function, or what they can be expected to accomplish. Furthermore, unlike the group examined in the Pilot Study, this group is not at all well known to the investigator.

Materials and Procedures

Four types of data were collected.

Table 3:
The Interdisciplinary Team: Stimulus Persons Chart

PS*	Identifier	Work Role	M/F?	Remarks
1	M_1	Medical Director	M	
2	M_2	Attending Physician-Outpatient Clinic	M	Recent graduate of residency.
3	M_3	Attending Physician-Outpatient Clinic	M	Recent graduate of residency.
4	M_4	Assistant Director, Head of Clinic Operations	M	Member of Department of Preventive Medicine
5	M_5	Attending Physician-Outpatient Clinic	M	Nominally a member of the Outpatient attending staff, but apparently this is not working well. More senior than M_3 or M_2.
6	M_6	Secretary	F	Keeps after everyone, especially with respect on the great body of paperwork. (For example, my questionnaires went to her for distribution to the Team.)
7	M_7	has been Coordinating Manager, now becoming Inpatient Counselor	M	Has not worked out in the administrative role- moving to a clinical one.
8	M_8	Physician's Assistant, Outpatient Clinics	F	De facto leader of the Outpatient Clinic group.

9	M_9	Supervising Social Worker, assigned to Outpatient Clinics	F	
10	M_{10}	Nurse Practitioner, Outpatient Clinics	F	
11	M_{11}	Pediatric Social Worker	F	
12	M_{12}	Physician's Assistant, Outpatient Clinics	M	Does blood testing on the Inpatient services with M_7.
13	M_{13}	Inpatient Social Worker	F	
14	M_{14}	Inpatient Caseworker	M	
15	M_{15}	Nurse Educator, Inpatient	F	
16	M_{16}	Nurse Educator, Inpatient	F	
17	M_{17}	Psychiatrist, Head of Social Work/Psychiatric	M	Nominally the leader of the SW/P division, but apparently not functioning effectively in that role.
18	M_{18}	Part-time Phlebotomist, Outpatient Clinics; Part-time clerical worker	F	Is being fired.

* PS indicates the plot symbol used for this individual in the output from the MDS, ADDTREE, and EXTREE algorithms.

The First Questionnaire

First, a two-part questionnaire was created, similar to those employed in the Jones and Young study and the Pilot Study. Part I asked the subjects to indicate "the number on a nine-point scale (from 1= don't know at all to 9= know very well) that best expresses how well you feel you know the person in question", for each of the 18 stimulus persons. Part II presented all of the possible pairings of the stimulus individuals in a "Ross ordering" (Ross, 1934), with the following instructions:

> "In this section, you are asked to rate the similarity of pairs of the X Team[16] staff, taking into account whatever characteristics you think are relevant. Please circle the number on the nine-point scale (from 1=extremely similar to 9= extremely dissimilar) which best represents the degree of similarity between the two."

Thirteen of the eighteen members of the Team completed the First Questionnaire. The similarity ratings collected from these subjects were then averaged and fitted to the

[16] The full title of the Team has been omitted to protect the privacy of the subjects.

ALSCAL, ADDTREE, and EXTREE models, as described in the Pilot Study.

The Organizational Chart

Second, the Medical Director of the Team was asked to draw a conventional organizational chart for the Team. This chart was be taken to be a "theory about the task structure" of the Team. In other words, it was taken to be a theory about the assignments of roles and tasks that had been established in order to carry out the overall task of the Team.

The Second Questionnaire

Third, the investigator compared the information provided by these two sources and developed a series of hypotheses about the task and sentience structures in the Team, using the sort of reasoning that was described above in the Pilot study. A focus of special interest was be the degree to which the subdivisions within the organization that are outlined in the organizational chart are, or are not, reflected in the analyses of the similarity data. In addition, a number of "distractor" hypotheses were formulated. These were statements that were not supported by the analysis of the similarity data.

Both types of statements were included in a second questionnaire which was given to the subjects with the following instructions:

> "In this questionnaire, you will be presented with a number of statements which express hypotheses regarding the structure of relationships within the X Team. Some of them may well be true, but some of them are very likely to be false. Most of these hypotheses concern the degree to which the four major organizational sub-units into which the X Team is divided (Outpatient, Inpatient, Social Work/Psychiatric, and Administrative) are "psychologically real" for the members of the Team. For each statement, you are asked to indicate whether it seems true or not, based on your knowledge of the Team. For each statement, please circle the number on the nine-point scale (from 1=absolutely false to 9=absolutely true) which best expresses your evaluation of the truth of the statement."

Twelve subjects completed the Second Questionnaire. Their responses were then collated and analyzed. To the degree that the inferences drawn by the investigator on the basis of the similarity data and the organizational chart drawn by the Medical Director were correct, it was predicted

that the subjects would tend to agree with the statements that reflect those hypotheses and disagree with those that do not (i.e. the "distractors").

The Interviews

Finally, all of the subjects who responded to the Second Questionnaire were invited to participate in a semi-structured interview conducted by the investigator. Ten agreed to do so. This interview began with a request to the individual to describe their role on the Team and then asked them to comment on each of the statements contained in the second questionnaire. The data from these interviews was used to illuminate and expand upon the "agreement data" described above.

Results of the Second Study

As discussed above, the data collection in this study fell into four stages. In the first, which was essentially a replication of the Pilot Study, similarity ratings were gathered, averaged, and the resulting averaged similarity matrix was then fitted to the multidimensional scaling, additive tree, and extended tree models.

Second, the leader of the Team was asked to draw a conventional organizational chart of the Team.

Third, these two representations were compared and a series of inferences about the functioning of the group was developed. Some of these inferences were included in a second questionnaire which was given to the subjects, who were asked to indicate whether they agreed with the statement or not.

Finally, subjects were asked to volunteer for a brief interview, in which they were asked to describe their roles and to add additional comments on each of the statements. Taken together, the interview and second questionnaire data provide evidence for the validity of the inferences drawn from the comparison of the similarity data and the organization chart.

The discussion below will deal with each of these phases, in turn.

The Results of the First Questionnaire

All eighteen persons who were members of the X Team at the time of the study were asked to complete the first questionnaire. Responses were received from thirteen

subjects. The triangular matrix of averaged similarity ratings is given as Table 4. This data was then fitted to the multidimensional scaling (ALSCAL), additive tree (ADDTREE), and extended tree (EXTREE) models. Table 5 presents a comparison of the variance-accounted-for and stress for all of these analyses.

Multidimensional Scaling

The data was fitted to three, four, and five dimensional spatial models. As indicated in Table 5, the improvements in variance-accounted-for and stress realized by the use of additional dimensions appear to be relatively modest, and only the three dimensional configuration will be discussed here.

Figures 9 and 10 present two views of the three dimensional configuration that resulted from the fitting of the similarity data. In contrast with the "pyramidal" arrangement that appeared in the analysis of the Pilot Study data, this configuration seems more like a sphere, with the leader (M_1) at one "pole" of an axis defined by dimension 2.

Though even this three-dimensional configuration is difficult to visualize, several observations can be made. First, if we take the leader's position to be the top of the

Table 4:
Averaged Dissimilarity Ratings for the Personnel of the Interdisciplinary Team

Individual																
R_2	5.5															
R_3	4.5	4.5														
R_4	3.7	3.7	4.6													
R_5	6.8	6.6	6.1	6.5												
R_6	4.8	6.8	5.2	5.8	7.2											
R_7	6.6	6.2	5.7	6.2	6.3	5.1										
R_8	6.0	4.8	5.0	5.5	6.5	4.5	4.9									
R_9	5.5	5.3	4.7	5.0	7.0	4.8	5.5	3.8								
R_{10}	4.0	4.8	4.4	4.5	6.5	5.2	6.8	4.2	5.2							
R_{11}	7.2	5.8	6.3	6.2	6.6	6.8	5.9	5.8	3.8	6.0						
R_{12}	7.3	5.5	5.9	6.0	6.2	6.0	3.3	6.2	5.2	7.2	5.6					
R_{13}	6.5	6.5	6.5	7.0	7.2	6.1	6.8	6.3	6.4	6.3	6.0	6.3				
R_{14}	6.8	6.9	7.1	6.7	5.9	4.9	5.9	6.8	5.5	6.4	6.4	5.8	3.1			
R_{15}	6.9	6.8	6.8	7.1	7.1	7.0	4.8	6.3	6.5	6.4	6.5	5.5	6.2	6.8		
R_{16}	7.1	7.2	7.2	6.9	6.9	7.2	5.1	6.2	6.1	6.0	6.1	5.5	6.8	5.9	2.8	
R_{17}	6.3	6.1	6.5	6.4	6.4	7.1	6.1	6.5	4.7	6.2	6.3	4.9	6.3	5.8	6.2	6.6
R_{18}	8.2	8.4	8.3	8.2	8.6	8.5	7.7	8.1	8.3	8.3	8.4	7.7	6.7	8.0	6.4	7.3	8.2
Individual	R_1	R_2	R_3	R_4	R_5	R_6	R_7	R_8	R_9	R_{10}	R_{11}	R_{12}	R_{13}	R_{14}	R_{15}	R_{16}	R_{17}

Table 5:
Variance-Accounted-For and Stress: Multidimensional Scaling, Additive Tree, and Extended Tree Models: The Interdisciplinary Team Data

Model	Variance Accounted For	Stress	
		Formula #1	Formula #2
Multidimensional Scaling:			
3-dimensions	64.6%	.220	
4-dimensions	66.4%	.176	
5-dimensions	69.8%	.142	
Additive Tree:			
Without Specified Initial Configuration	80.75%	.075	.328
With Organization Chart as an Initial Configuration	50.23%	.105	.539
Extended Tree:			
Without Specified Initial Configuration	88.63%	.046	.240
With Organizational Chart as an Initial Configuration	78.60%	.071	.344

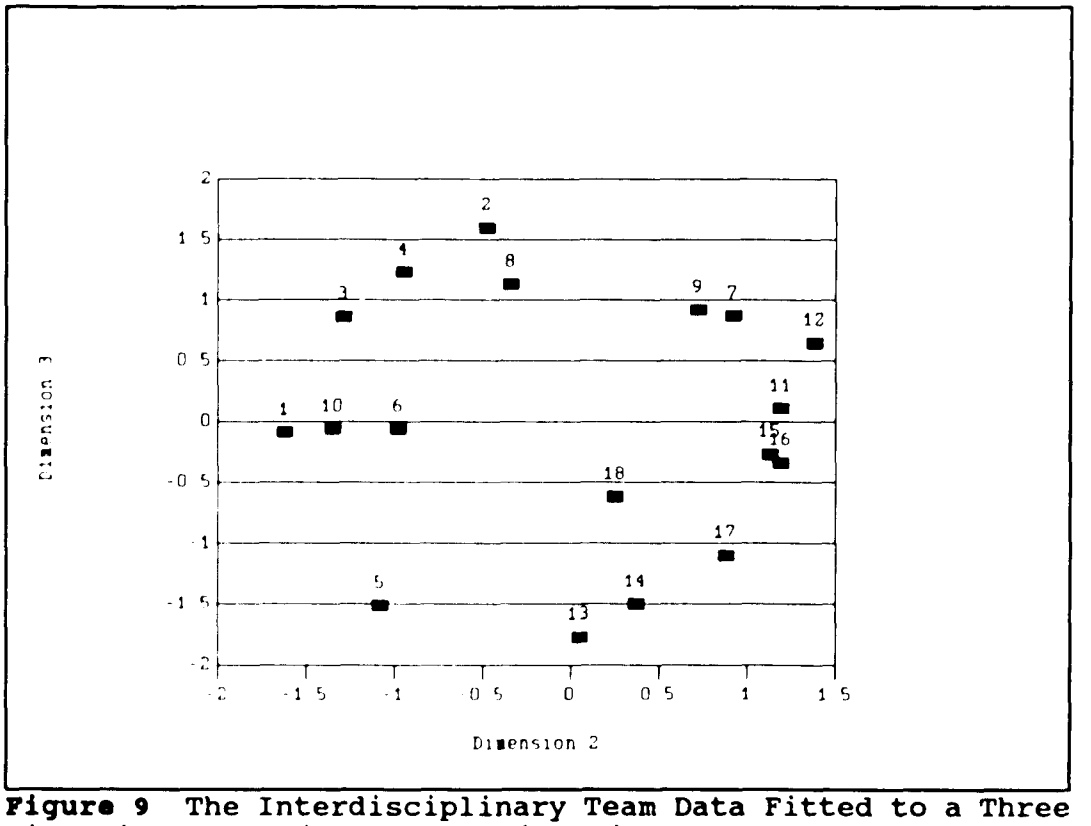

Figure 9 The Interdisciplinary Team Data Fitted to a Three Dimensional Spatial Model (Dimensions 2 and 3)

sphere, it seems that the upper hemisphere contains most of the physicians (M_1, M_2, M_3, M_4, and M_5). The only exception is M_{18}. Three physicians (M_2, M_3, and M_4) cluster together on one side of the sphere, while another (M_5) is off at some distance on the other side. Sharing the upper hemisphere with the physicians are two of the Outpatient staff (M_8 and M_{10}) and the Secretary. The lower half of the sphere contains the Social Work staff (M_9 and M_{11} on one side, and M_{13} and M_{14} on the other).

Second, it would appear that there are, within this

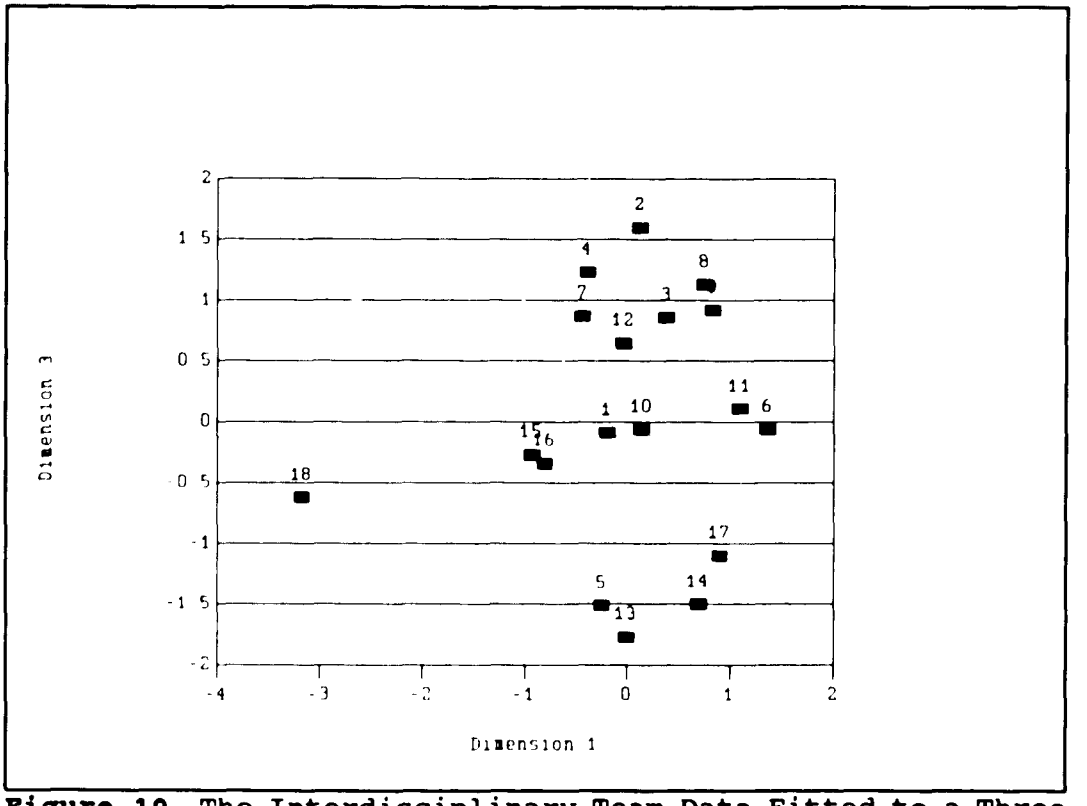

Figure 10 The Interdisciplinary Team Data Fitted to a Three Dimensional Spatial Model (Dimensions 1 and 3)

group, a number of tightly bound pairs (M_{13} and M_{14}, M_{15} and M_{16}, M_7 and M_{12}). Overall, however, the multidimensional scaling results do not appear to be very illuminating, at least taken by themselves.

ADDTREE and EXTREE

Figures 11 and 12 present the ADDTREE and EXTREE output configurations that were obtained from the analysis of the

Figure 11:
The Interdisciplinary Team Data Fitted to an
Additive Tree Model (ADDTREE Algorithm)

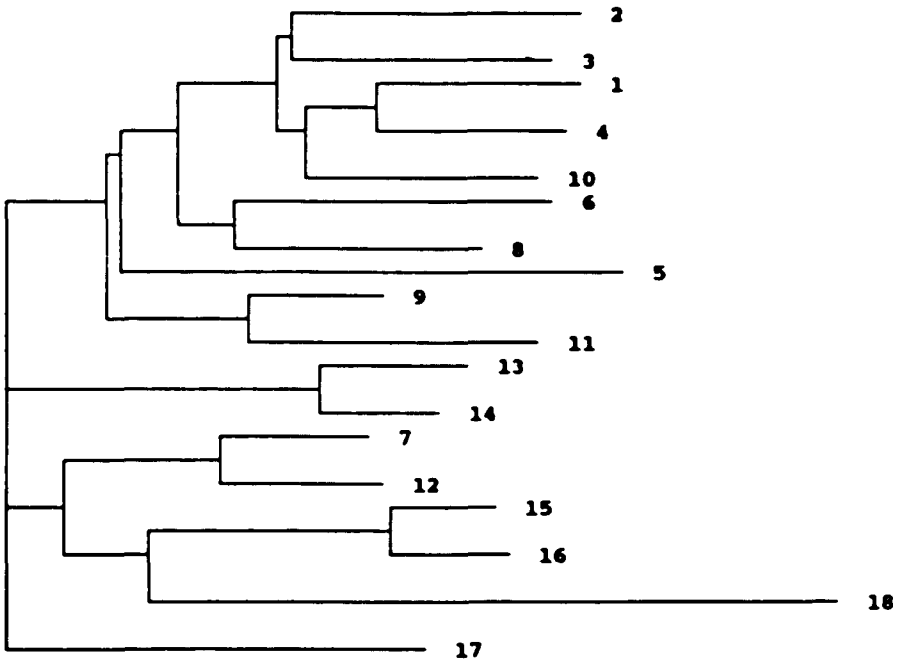

Figure 12:
The Interdisciplinary Team Data Fitted to an
Extended Tree Model (EXTREE Algorithm)

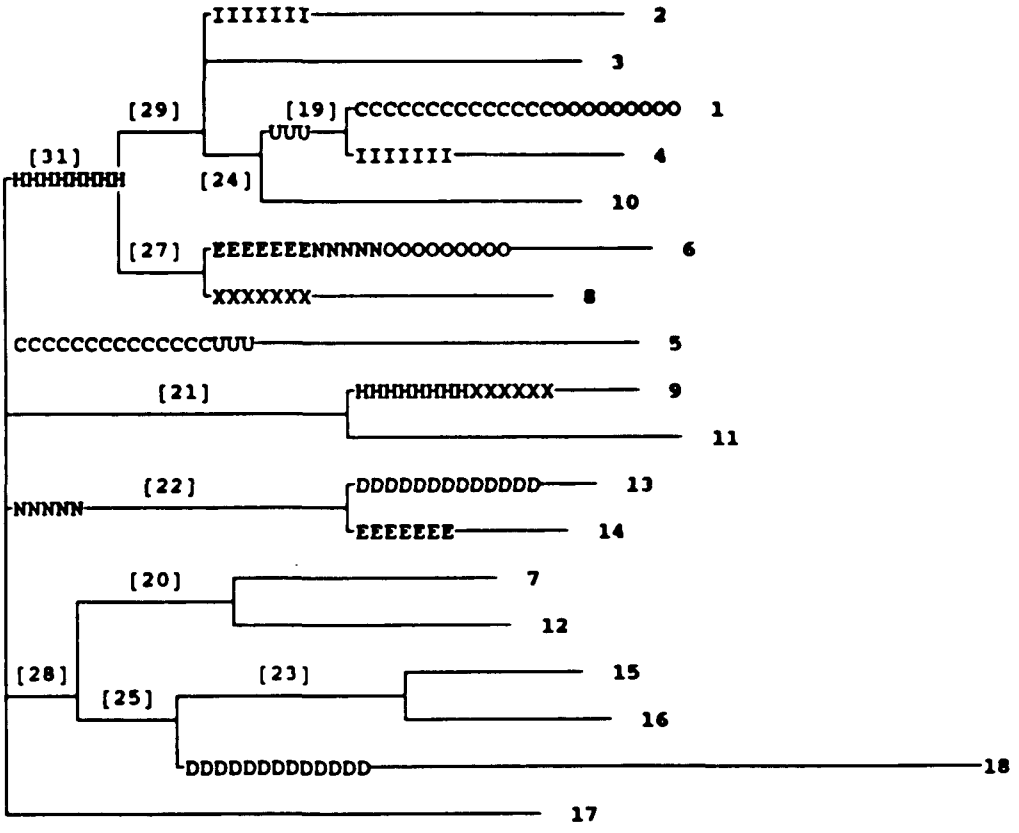

Cross-Cutting Categories:

Identifier:	Members:
C	[1, 5]
D	[13, 18]
E	[14, 16]
H	[9, 1, 2, 3, 4, 6, 8, 10]
I	[4, 2]
N	[6, 13, 14]
O	[1, 6]
U	[1, 4, 5]
X	[9, 8]

Table 6:
Extensions and Weights for Clusters:
EXTREE Analysis of the Interdisciplinary Team Data

Cluster:*	Extension or Weight Assigned	
	Extension/Weight	Z-score
Tree Branches:		
19	0.46	-0.51
20	0.78	0.42
21	1.00	1.05
22	1.69	3.05
23	1.14	1.46
24	0.27	-1.06
25	0.49	-0.42
27	0.41	-0.66
28	0.33	-0.89
29	0.42	-0.63
31	0.59	-0.13
Cross-Cutting Clusters:		
C	0.97	0.97
D	0.88	0.71
E	0.50	-0.40
H	0.59	-0.13
I	0.59	-0.13
N	0.37	-0.77
O	0.62	-0.05
U	0.23	-1.18
X	0.40	-0.69

* Clusters are labelled in Figure 13

similarity ratings[17]. This discussion will focus on the EXTREE results.

Four general observations can be made on the basis of the general appearance of the configuration.

First, the configuration includes one fairly large cluster, which includes seven of the eighteen stimulus individuals (stimulus persons M_2, M_3, M_1, M_4, M_{10}, M_6 and M_8), which will referred to as the "main cluster." There is another cluster of five individuals (containing stimulus persons M_7, M_{12}, M_{15}, M_{16}, and M_{18}), but, at first inspection, its distinctiveness (as indicated by the length of the branch labelled [28] which separates it off from the rest of the group) appears to be somewhat less than that of the other clusters. Furthermore, the extensions of its subclusters (the cluster labelled [20] containing individuals M_7 and M_{12} and the cluster labelled [23], containing individuals M_{15} and M_{16}) are rather large in comparison to the extension of the node that separates it from the whole group, suggesting that the distinctions within this putative subcluster may be more significant than the distinction between the cluster as a whole and the rest

[17] The bracketed numbers in Figure 13 serve to label the extensions for each of the branches in the tree structure.

of the group.

All of this suggests that it may be a less meaningful subgroup than the main cluster. This initial impression is supported by a close examination of the extensions assigned to each of the clusters in the EXTREE output, which can be thought of as reflecting the salience of the boundary that separates the subcluster from the rest of the group. Table 6 gives the extensions assigned to each of the clusters in the tree structure and to the cross-cutting categories. The mean extension length of all of the clusters is .48 with a standard deviation of .40, and these statistics were used to calculate z-scores (based on the set of the extensions of the clusters, as given in Table 6) for each extension. If we of compare the extensions of the nodes for the main cluster (labelled [31]) and the 5-member cluster containing individuals M_7, M_{12}, M_{15}, M_{16}, and M_{18} (labelled [28]) we find that the main cluster's node is slightly larger than average, while the node for the other cluster is slightly below average. The impression that the extensions of the subclusters of this 5-member cluster are more salient than the cluster as a whole is also supported by the extension data. If we calculate the ratios of the extensions of the sub-clusters to the extensions of the whole clusters, we find that, in the main clusters, these ratios are less than one, but that in the 5-member cluster they are larger than

one. It would appear, therefore, that the boundary that separates the 5-member cluster from the rest of the group than is less salient than the boundaries that separate its sub-clusters from each other. This suggests that one should be cautious in attributing any degree of "psychological reality" to this cluster.

The main cluster is also noteworthy in that the Medical Director (M_1) is located in its innermost layer, at the center of the cluster. This placement is reminiscent of some of the clusters that contained the Vice-CEO's (especially the Vice-CEO for General Administration) in the Pilot Study results.

A second finding is the prevalence, outside the main cluster, of pairs of individuals who are very strongly bound to each other but have relatively weak connections to others in the group. Four such pairs, comprising eight of the eleven individuals outside the main cluster, can be identified: M_9 and M_{11} (the cluster labelled [21]); M_{13} and M_{14} (the cluster labelled [22]; M_7 and M_{12} (the cluster labelled [20]); and M_{15} and M_{16} (the cluster labelled [23]). These pairs might be called "satellites." If we were to disregard the 5-member cluster, this pattern becomes even more striking. Furthermore, the length of the extensions for these four clusters are among the largest. Graph 4

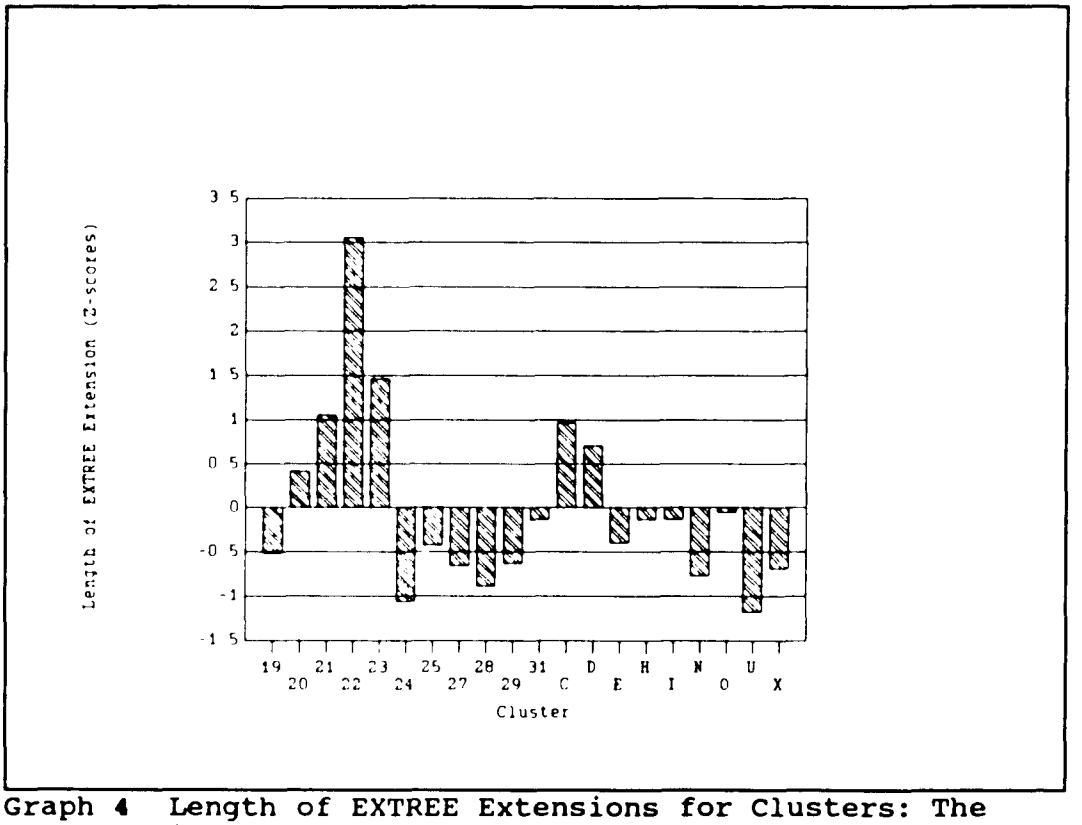

Graph 4 Length of EXTREE Extensions for Clusters: The Interdisciplinary Team Data

presents the z-score data from Table 6, and shows that all four of the above mentioned clusters are among the six clusters with extension lengths greater than +.50. Three of them have extension lengths greater than +1.0. This would seem to indicate that these pairings have great salience for the members of the group.

Third, two individuals (M_5 and M_{17}) are set off by themselves, and a third (M_{18}) is located within a rather weakly marked cluster structure and has been given a very long extension. These are signs that are somewhat

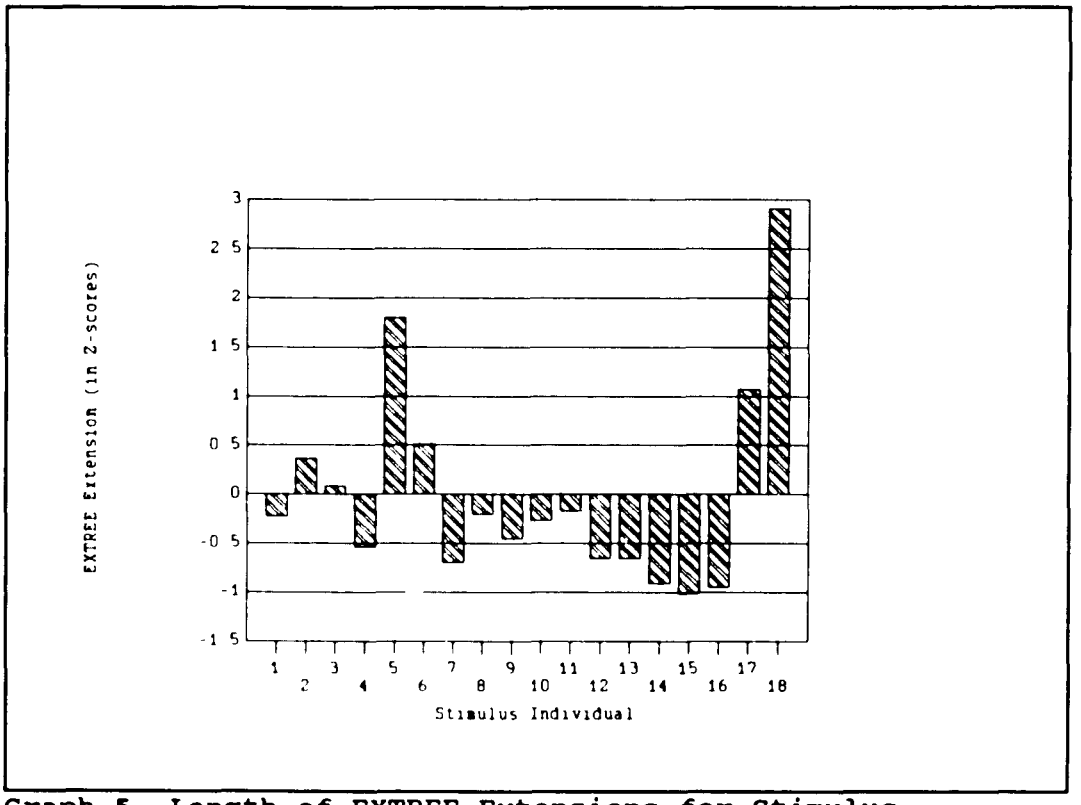

Graph 5 Length of EXTREE Extensions for Stimulus Individuals: The Interdisciplinary Team Data

suggestive of "outcast" status, as discussed in the Pilot Study. Graph 5 displays the extensions for all of the stimulus individuals (in z-scores). Using the method that was used in the Pilot Study, it would appear that the extensions for these three individuals are much larger than average. They are in fact, the only individuals with extensions whose z-scores are greater than +1.0. Thus, they seem to fulfill the quantitative and qualitative criteria for "outcast" status.

Finally, it is interesting that six of the nine cross-

cutting clusters that the EXTREE algorithm superimposed on the original ADDTREE tree structure serve to make linkages between the main cluster and individuals or pairs outside the main cluster. Furthermore, two of the other three cross-cutting clusters serve to link individuals in outer layers of the main cluster to the innermost layer, where the leader (the Medical Director) is located. Thus, the majority of the cross-cutting clusters serve to tie individuals back to the main cluster or to the leadership, rather than to link these "satellites" together.

Overall, the picture that emerges from the data is of a group which consists of three elements:

1.) a relatively large main body, centered on the group leadership;

2.) four pairs of individuals who are very tightly bound to each other but only loosely linked to the rest of the group, generally through connections back to some part of the main cluster; and

3.) three isolated individuals (possibly "outcasts"- individuals who are having difficulty taking up their role).

Two questions then need to be considered. First, is this picture accurate? Second, if it is accurate, what is its relation to the task structure for this group? The first step in answering these questions was to request an organization chart from the Medical Director.

The Organizational Chart

No formal organization chart for the Team was available at the time of the study and so the investigator asked the Director to draw one. The sequence in which the sections of the chart were drawn and the comments with which he accompanied it are interesting, and worth including. A simplified version of the Medical Director's chart is given as Figure 13.

The Director (M_1) began with his own position, placing himself at the top center, the conventional place for the leader.

He then added, below and to the left, the Assistant Director (M_4). At first, he had been inclined to place the Assistant Director directly below, but stopped himself, remarking that the Assistant Director's responsibilities did not encompass the whole Team, but only the Outpatient Clinic side. He then added the notation "Clinic Operations" in the

Figure 13:
The Medical Director's Organizational Chart for the Interdisciplinary Team

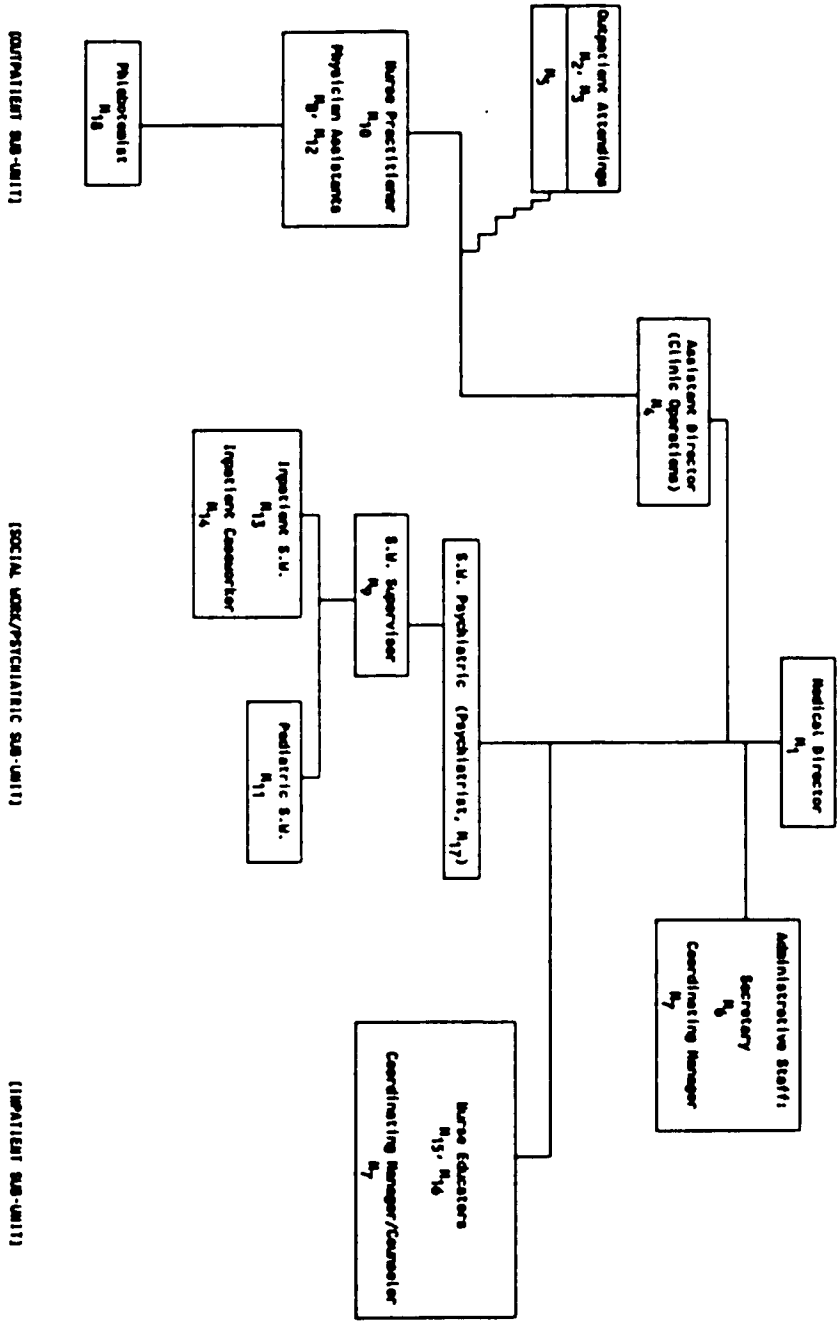

Assistant Director's box.

Next, he drew a line (with two right-angle bends in it) from the Assistant Director to a box in which he placed the two Physician's Assistants (M_7 and M_8) and the Nurse Practitioner (M_{10}). The box represented the core of the Outpatient Clinic staff. One of the Physician's Assistants (M_8) was identified as being the person that he saw as having the closest thing to a leadership role in this sub-group.

The Director then drew a line extending out of the line connecting the Assistant Director and the Outpatient staff, and added a box which, at first, contained only two of the Outpatient Attending staff (M_2 and M_3). Later, he added an additional section to this box, and included another attending (M_5), commenting that though it had originally been hoped that he would coordinate Inpatient activities, this "had not worked out." In describing the role of the attending physicians on the Team, the Director remarked that they were "resources, problem solvers, provided higher level administration."

The Director then moved on to the second major sector of the Team, which he labeled as "S.W. [Social Work] Psychiatric." He drew a curved line down to M_9, who was

labelled as the S.W. Supervisor. Below her, he split the group into two sub-groups. One, containing M_{13} and M_{14}, was identified as the social workers assigned to the hospital's Inpatient Services. The other was the social worker assigned to the Pediatrics service (Inpatient and Outpatient), M_{11}.

The Director then returned to the Social Work Psychiatric box and added in its putative leader, M_{17}, commenting that he was an MD who had been expected to work on staff development, but that "it hadn't worked out." He described the function of the Social Work Psychiatric unit as providing support services to both the Inpatient and Outpatient sides of the Team. Furthermore, all of the staff (except, perhaps, the Psychiatrist) have direct service responsibilities through either the Inpatient or Outpatient sides of the Team.

The Director then added the third major component of the Team, the Inpatient. In this box, he placed the two Nurse Educators (M_{15} and M_{16}) and the Coordinating Manager (M_7) noting that the latter was in transition from an administrative role to a service delivery role (as a Counselor).

The final component added was the administrative

component, which he drew as an offshoot of his own position. In this box, he placed the Coordinating Manager (M_7), though he repeated that this person was in transition out of this role, and the Secretary (M_6). In describing the Secretary's role, he commented that she had very heavy clerical responsibilities and provided support to all of the Team components.

The last step in drawing this map of the Team was to add in the Phlebotomist (M_{18}), who was given two locations. First, she was appended to the Outpatient group, identified as part-time and it was noted that one of the Physician's Assistants (M_8) was responsible for monitoring her time and attendance. Then, he added her to the administrative group, indicating that she had part-time clerical responsibilities as well. He also commented that she was "on the outside" and that she was "being fired."

The Two Representations Compared

At this point, there were two different representations of the Team, derived from very different sources. One is the picture derived from the reconstruction of the "maps-in-the-head" from which we hypothesize that the similarity data have been drawn. The other is the organization chart and the comments made by the Director as he drew it. In this

section, we will discuss the ways in which these two types of data may be woven together to develop the specific and testable inferences about the Team which were incorporated into the Second Questionnaire. We will begin with the organization chart and the concept of the task structure for the Team that it embodies.

According to the chart, the Team consists of four major components. Two of these components (Inpatient and Outpatient) provide direct clinical services to patients. They appear to be operating units whose throughput is patients. A third (Social Work Psychiatric) does not appear to have been intended as a way to provide direct care, although it houses staff who do provide care, but rather what Miller and Rice called a "disciplinary base", a "home" for staff with a common disciplinary affiliation whose direct service work is done through the other units. The fourth (the administrative unit) is directed toward the management and support of the Team and to the regulation of its relationships with the environment. Thus, the Team appears to contain three somewhat different types of sub-units and, most importantly, many members of the Team have affiliations to more than one sub-unit.

Now, suppose we were to superimpose the representation provided by the analysis of the similarity data onto the

representation provided by the organization chart- how do they correspond to each other?

First, let us consider the degree to which the four major organizational sub-units listed above appear to be "psychologically real"[18]:

1. The Administrative Sub-Unit

The Medical Director and the Secretary are both placed in the main cluster (though in different areas) and are linked by being the sole members of a cross-cutting cluster ("O") with a fairly substantial extension (.62). Thus, it would appear that these is some "psychological reality" to this part of the Administrative Sub-Unit. There is no evidence that the third member of the Sub-Unit (the Coordinating Manager) is thought of as being a "real member" of the administrative group, an impression that the Medical Director's comment that the Coordinating Manager was in transition out of this role would seem to confirm.

---

[18] "Psychologically real" will be used to mean the degree to which membership in a group influences perceived similarity.

2. The Inpatient Sub-Unit

The three staff identified by the Medical Director as being a part of this sub-unit are all found in the five-member cluster discussed above, but this cluster also includes two individuals (one of the Physician's Assistants and the Phlebotomist) whose primary assignments are to other sub-units, and so it is difficult to interpret this as an Inpatient group. What is more striking about this sub-unit is the very marked degree to which the two Nurse Educators are distinguished from the rest of the group- this group seems to have a high degree of salience. The formal Inpatient assignment mentioned by the Medical Director for M_5 appears to have no "psychological reality."

3. The Outpatient Sub-Unit

For the most part, the members of this group are found within the main cluster. The exceptions are M_5, one of the Physician's Assistants, and the Phlebotomist.

4. The Social Work/Psychiatric Sub-Unit

> The five members of this group are found in two of "satellite" pairings (one consisting of the two Inpatient Social Work staff, the other consisting of the Pediatric Social Worker and the Supervising Social Worker) and in one of the "isolates" (the Psychiatrist). Thus, it would appear that though some of the components of this unit have great salience, the sub-unit itself, as a whole, has no palpable degree of "psychological reality." Rather than being bound to each other, it would appear that each of the two satellites is connected back to the main cluster (through cross-cutting clusters "H", "X", "D", "N", and "E").

Second, let us consider the question of the "outcasts." As discussed above, the analysis of the similarity data suggested that individuals might be having difficulty taking up their role in the group: 1) M_5; 2) the Psychiatrist; and the Phlebotomist. Though the formal organization chart was concerned with showing where they "should be", there were several indications, in the comments made by the Director and in the way in which he drew the chart, about where they actually were:

1. M_5 was added to the group of Outpatient Attendings as an afterthought;

2. The Psychiatrist was initially left out of the group of which he is, on paper, the leader and the Director said that "it hadn't worked out"; and

3. The Phlebotomist is "on the outside", her time and attendance requires special monitoring, and she is "being fired."

On balance, the data from the chart support the identification of these individuals as "outcasts."

EXTREE with the Organizational Chart as an Initial Configuration

The discussion above would seem to indicate that there are some marked discrepancies between the two representations of the organization drawn from the organization chart and the analysis of the similarity data. One of the interesting properties of the ADDTREE and EXTREE procedures is that they provides for a very direct test of this impression. The task structure (as embodied in the organization chart) can be translated into a tree structure which can then be used as an initial configuration for the

ADDTREE and EXTREE algorithms. Instead of beginning with a random configuration, the computer begins with a configuration that reflects the sentience pattern that one would expect if the judged dissimilarities were entirely determined by the task structure. To the extent that the actual similarity data depart from this expectation, this initial configuration will be adjusted in the course of the iterative algorithm. In a manner of speaking, one begins with a state that corresponds to hypothesis is that the similarity ratings are entirely determined by the functional divisions specified by the task structure. If that hypothesis is correct, the configuration would fit the similarity data well, and would not be adjusted in the course of the iterative process of the algorithm. If it incorrect, the algorithm would, inevitably, move away from the initial configuration as it attempts to improve the ability of the configuration to explain the similarity data.

A tree structure derived directly from the Medical Director's organization chart (see Figure 13) was used as an initial configuration for both ADDTREE and EXTREE. Figure 14 presents the EXTREE output configuration that resulted, which accounted for 78.60% of the total variance. In its general features, it strongly resembles the EXTREE output obtained from a random starting point. The major difference is that the main cluster does not appear immediately, though

Figure 14:
The Interdisciplinary Team Data Fitted to the
Extended Tree Model Using the
Organizational Chart as an Initial Configuration

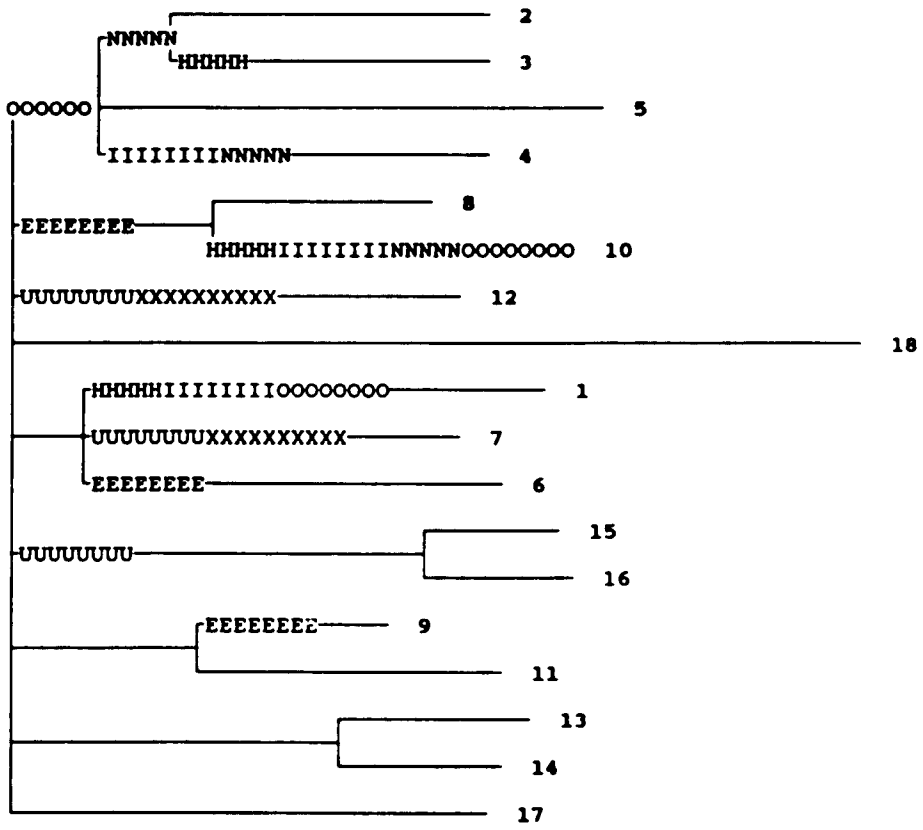

Cross-Cutting Categories:

Label:	Members:
E	[6, 8, 10, 9]
H	[3, 1, 10]
I	[4, 10, 1]
N	[4, 10, 2, 3]
O	[10, 2, 3, 4, 5, 1]
U	[12, 7, 15, 16]
X	[7, 12]

it does appear in a modified form. The individuals who make up the main cluster in the first analysis are found in three separate clusters in this analysis but the cross-cutting clusters "H", "I", "N", "O", and "E" serve to link them together in much the same way.

Furthermore, the configurations that were obtained using the random starting points for both EXTREE and ADDTREE accounted for more variance than the configurations that were obtained using the organizational chart as a starting point (see Table 5, page 121). On the whole, the validity of the picture of the organization presented by the original configurations seems to be supported. This was tested further through the Second Questionnaire.

The Construction of the Second Questionnaire

On the basis of the foregoing analysis, sixteen statements about the Team were prepared for testing in the Second Questionnaire. These statements can be grouped into three categories:

Group I: Statements which are not supported at all by either the analyses of the similarity data or the organization chart (#9, #10, #12, and #14). These statements were included to

provide a check on the suggestibility of the subjects.

Group II: Statements which are either weakly or equivocally supported by the similarity data (#2, #4, and #11) or supported by the organization chart but not by the similarity data (#8):

#2. **"The Outpatient staff are a tightly knit group."**

The existence of the main cluster, which includes most of the Outpatient staff, may be taken as weak support for this statement but, for the most part, the data provides no evidence either for or against this statement.

#4. **"The Phlebotomist is closely associated with the two Nurse Educators."**

As noted above, though the Phlebotomist is joined to the two Nurse Educators in a cluster, that cluster has an extension of only average size, and there are

grounds for suspecting that it may be spurious.

#8. **"The Psychiatrist and the Supervising Social Worker are closely associated."**

This statement is consistent with the structure described by the organization chart, but finds no support in the similarity data.

#11. **"The Secretary serves as a central connecting link for the Team, tying the many subunits together."**

The indication for this statement comes from the Medical Director's description, from which it would appear that this is the intended role for the Secretary, and the fact that the Secretary is a member of four of the nine cross-cutting categories.

Group III: Statements which are strongly supported by the similarity data:

#1. **"M_5 and M_1 are closely associated."**

The indication for this statement is the cross-cutting category "C" in the EXTREE analysis, which has an extension of .97.

#3a,b. **"The Outpatient staff group contains two distinct sub-groups: a.) the three attending physicians; and b.) the Social Worker and one of the Physician Assistants."**

The indications for the first part of this statement are the large sub-cluster (extension= .42) within the main cluster that contains the Outpatient attending staff and the cross-cutting category "I" (extension= .59), which links two of the attending staff. The indications for the second part are the cross-cutting cluster "X" (extension= .40). It should be noted that the extensions for all of these indications are relatively small, suggesting that the suggested groupings are not highly salient.

#5a,5b. **"The Social Work/Psychiatric staff group contains two distinct sub-groups: a.) The Supervising Social Worker and the Pediatric Social Worker; and b.) The Inpatient Social Worker and the Inpatient Case Worker."**

The two sub-groups mentioned here correspond to the "9-11" (extension= 1.0) and the "M_{13}-M_{14}" (extension= 1.69) clusters. These are two of most highly distinctive clusters and so these statements can be seen as strongly supported by the similarity data.

#6. **"In general, the physicians on the Team have more in common with each other than they do with the non-physician staff of the units (e.g. Outpatient, Inpatient) to which they are assigned."**

Four of the six physicians are found in one branch of the main cluster. The remaining two are among the possible

"outcasts", but one of these is strongly linked (see #1 above) to the Medical Director. In general, the physicians appear to constitute a reasonably distinct sub-group within the Team.

#7. **"These two subunits of the Social Work/Psychiatric staff are each more closely associated with their places of work (Inpatient, Outpatient) than they are with their disciplinary base (Social Work/Psychiatric)."**

This statement was an attempt to directly express the impression that the Social Work Psychiatric Unit described by the Medical Director was not "psychologically real." The indications for it are primarily negative ones (i.e. that it does not appear as a distinct group) and the two cross cutting categories ("H" extension= .59 and "X" extension= .40) which link the Supervising Social Worker to the main cluster. In the absence of indications that the disciplinary base was not

"psychologically real", it seemed likely that the affiliations to the two service delivery units would predominate.

#13. **"The Supervising Social Worker works closely with the physician leadership of the Team."**

The indication for this statement is the cross-cutting clusters ("H", extension= .59) which links the Supervising Social Worker to the main cluster.

#15. **"The two Nurse Educators form a distinct sub-unit within the Team."**

The indication for this statement is the "M_{15}-M_{16}" cluster, which has an extension of 1.20. This is a relatively large extension, and so this statement can be considered to be strongly supported by the similarity data. It is also, of course, supported by the organizational chart.

#16. "One of the Physician's Assistants and the Coordinating Manager form a distinct sub-unit within the Team."

The indication for this statement is the "M_7-M_{12}" cluster, which has an extension of .80. This is also a relatively large extension, and so this statement can also be considered to be strongly supported by the similarity data. It is not supported by the organizational chart.

Direct statements about the status of the three possible "outcasts" were avoided, because it was thought they might be painful or damaging to the individuals involved.

The Results of the Second Questionnaire

The Second Questionnaire was distributed to all members of the Team. Twelve subjects completed and returned the Questionnaire. Ten of those subjects volunteered to be interviewed. Table 7 gives their responses to the sixteen statements in the Questionnaire, in both raw and within-subject z-scores. Graph 6 presents the averaged within-

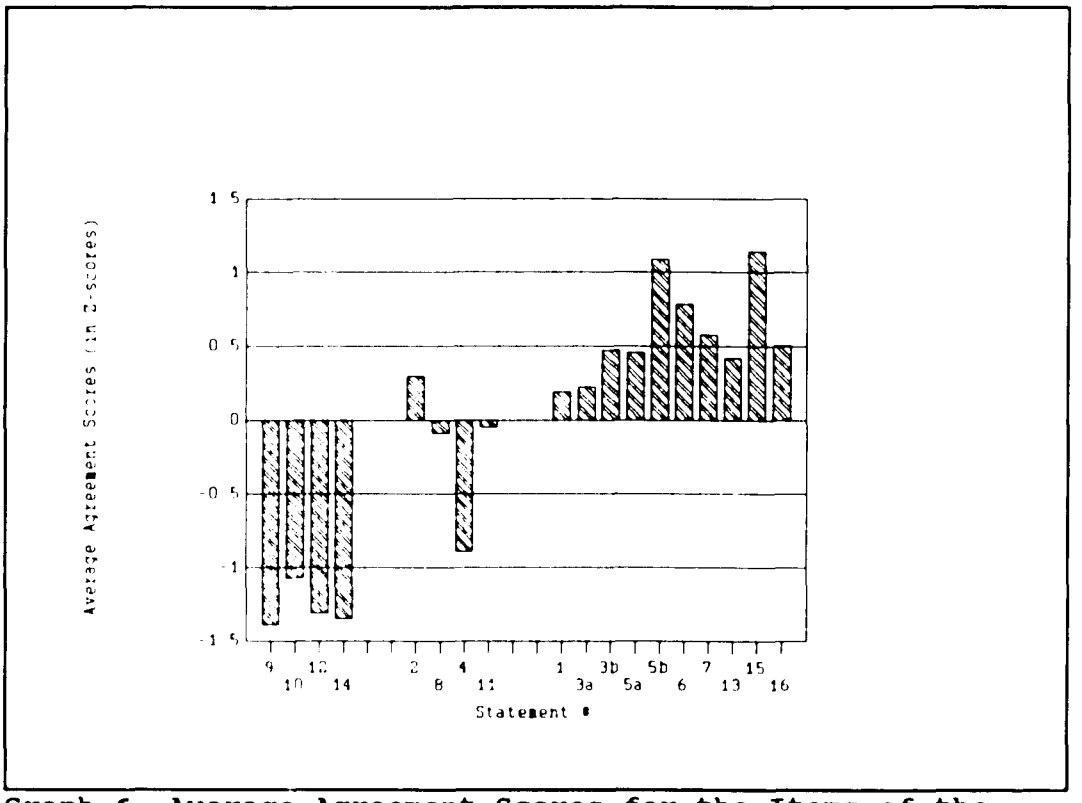

Graph 6 Average Agreement Scores for the Items of the Second Questionnaire

subject z-scores for the statements, which have been grouped as described above, and Graph 7 presents the averages for each group. At this coarse level of analysis, it would appear that there a distinct relationship between the degree to which the statements are supported by the findings of the similarity data and the degree to which they are endorsed by the respondents to the Second Questionnaire.

This impression can be tested more carefully by comparing the degree to which statements were endorsed by the subjects (agreement scores) with the strength of the

Table 7:
Responses to the Second Questionnaire

Statements	Individual Subject Responses												Averages	
	1	2	3	4	5	6	7	8	9	10	11	12		
	Raw Zsc	Raw Zsc	Raw Zsc	Raw Zsc	Raw Zsc	Raw Zsc	Raw Zsc	Raw Zsc	Raw Zsc	Raw Zsc	Raw Zsc	Raw Zsc	Zsc	Raw
1	7 0.33	8 0.82	8 1.09	2 -1.92	3 -1.30	4 -0.52	8 1.34	6 0.10	9 1.11	5 -0.18	8 0.98	7 0.42	0.19	6.25
2	9 0.88	8 0.82	7 0.67	3 -1.44	9 0.89	7 0.56	1 -0.97	6 0.10	5 -0.82	7 0.82	7 0.59	8 0.84	0.38	6.42
3a	9 0.88	8 0.82	3 -0.99	8 0.96	9 0.89	5 -0.14	3 -0.31	6 0.10	8 0.83	4 -0.57	5 -0.20	7 0.42	0.22	6.25
3b	9 0.88	5 -0.34	7 0.67	8 0.96	9 0.89	7 0.56	3 -0.31	3 -1.21	9 1.11	6 0.22	8 0.98	9 1.26	0.47	6.92
4	1 -1.34	1 -1.98	2 -1.41	6 0	4 -0.94	1 -0.52	1 -0.97	3 -1.21	1 -1.14	5 -0.18	6 0.20	3 -1.26	-0.89	3.08
5a	9 0.88	8 0.82	8 1.08	6 0	8 0.53	6 0.20	2 -0.64	9 1.40	2 -0.86	6 0.22	8 0.98	8 0.84	0.45	6.67
5b	9 0.88	8 0.82	8 1.08	8 0.96	9 0.89	9 1.29	8 1.34	9 1.40	9 1.11	8 1.01	9 1.38	8 0.84	1.08	8.50
6	9 0.88	8 0.82	8 1.08	9 1.44	6 -0.20	9 1.29	9 1.67	7 0.53	9 1.11	8 1.01	6 0.20	5 -0.42	0.78	7.75
7	9 0.88	6 0.96	7 0.67	8 0.96	7 0.16	7 0.56	5 0.33	8 0.97	9 1.11	7 0.82	6 0.20	7 0.42	0.58	7.17
8	1 -1.34	7 0.43	5 -0.16	8 0.96	7 0.16	4 -0.52	4 0.02	7 0.53	7 0.55	7 0.82	4 -0.59	2 -1.68	-0.08	5.25
9	1 -1.34	2 -1.51	2 -1.41	4 -0.96	2 -1.47	1 -1.41	1 -0.97	3 -1.21	1 -1.14	1 -1.76	2 -1.38	2 -1.68	-1.39	1.83
10	1 -1.34	3 -1.12	3 -0.99	2 -1.92	5 -0.57	2 -1.25	1 -0.97	3 -1.21	1 -1.14	2 -1.37	2 -1.38	7 0.42	-1.07	2.67
11	9 0.88	7 0.43	6 0.25	6 0	9 0.89	6 0.20	1 -0.97	4 -0.77	1 -1.14	7 0.82	2 -1.38	7 0.42	-0.05	5.42
12	1 -1.34	2 -1.51	2 -1.41	5 -0.48	2 -1.47	1 -1.41	1 -0.97	3 -1.21	1 -1.14	1 -1.76	2 -1.38	5 -1.26	-1.31	2.00
13	6 0.05	7 0.43	7 0.67	6 0	9 0.89	8 0.53	6 0.66	8 0.97	2 -0.86	7 0.82	5 -0.20	8 0.84	0.42	6.58
14	1 -1.34	2 -1.51	2 -1.41	5 -0.48	2 -1.47	1 -1.41	1 -0.97	3 -1.21	1 -1.14	1 -1.76	2 -1.38	2 -1.68	-1.35	1.92
15	9 0.88	9 1.21	8 1.08	7 0.48	9 0.89	9 1.29	9 1.67	9 1.40	9 1.11	9 1.41	9 1.38	8 0.84	1.14	8.67
16	5 -0.23	7 0.43	6 -0.58	7 0.48	9 0.89	7 0.93	7 1.01	7 0.53	7 0.55	7 0.82	8 0.98	7 0.42	0.56	6.92

indications for those statements. Table 8 presents the

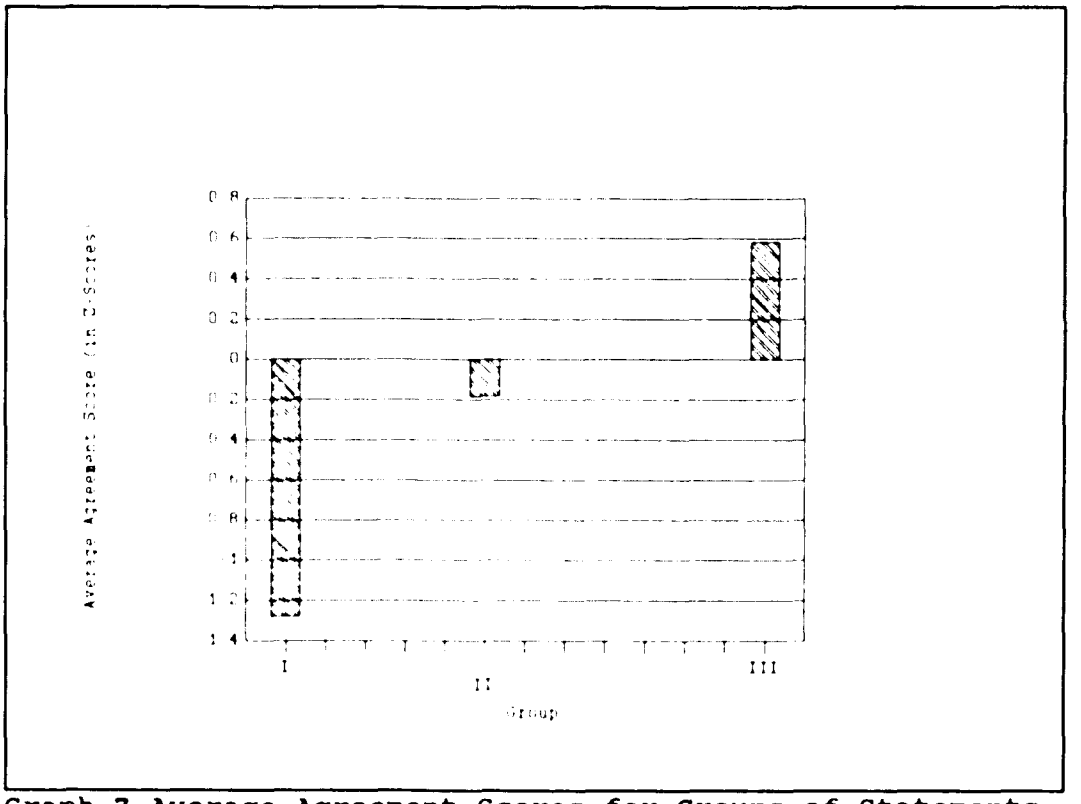

Graph 7 Average Agreement Scores for Groups of Statements

average agreement score (in raw and within-subject z-scores) for 14 of the 16 statements[19] in the second questionnaire and the strength of the indications (as measured by the length of the EXTREE extension that supports the statement) for each statement. Figure 15 is a plot of average agreement against the strength of the indication. A roughly linear relationship appears to hold between the two. In

[19] Statements #2 and #6 are not included as they are inferences based on an absence of information and, therefore, there are no quantifiable indications in the similarity data.

Table 8:
Strength of Indication and Agreement Scores

Statement	Average Agreement Scores:		Strength of Indication
	Z-scores	Raw Score	
1	0.19	6.25	0.97
3a(1)*	0.22	6.41	0.42
3a(2)*	0.22	6.41	0.59
3b	0.47	6.91	0.40
4	-0.89	3.09	0.49
5a	0.45	6.67	1.00
5b	1.08	8.50	1.69
7	0.58	7.17	0.59
8	-0.09	5.25	0
9	-1.39	1.83	0
10	-1.07	2.67	0
11(1)*	-0.05	5.42	0.50
11(2)*	-0.05	5.42	0.37
11(3)*	-0.05	5.42	0.62
12	-1.31	2.00	0
13	0.42	6.58	0.59
14	-1.35	1.91	0
15	1.14	8.67	1.20
16	0.50	6.91	0.80

* Some statements have several different indicators

order to determine this relationship more precisely, a correlation coefficient (Pearson's r) was calculated for seventeen of the 19 statements in the Second Questionnaire. (Statements #12 and #14 were dropped as they are duplicates of Statement #9.) A correlation coefficient of .763 was obtained for the strength of indications and raw agreement scores, and a coefficient of .767 was obtained between agreement z-scores and strength of indications (for both, $p<.01$). It might be argued, however, that this relationship has been skewed by the inclusion of several items (e.g. Statement #9) which had no support in the data and were, therefore assigned a strength of indication of 0 in Table 8. One could imagine a version of the Second Questionnaire that included a large number of items such as "the moon is made of green cheese"- statements which have no support in the data and which would also be given very low agreement scores. By including enough such items, the correlation coefficient could be raised to virtually any desired level. Therefore, a second correlation was computed, in which all of the items with a strength of indication of 0 (Statements #8, #9, #10, #12, and #14) were discarded. The resulting correlation coefficients were .668 for z-scores and .651 for raw scores, both with $p<.01$. It appears that there is a significant correlation between strength of indication and agreement scores.

Two of the exceptions (#8 and #4) to this relationship are interesting. Statement #8 appears to be somewhat more

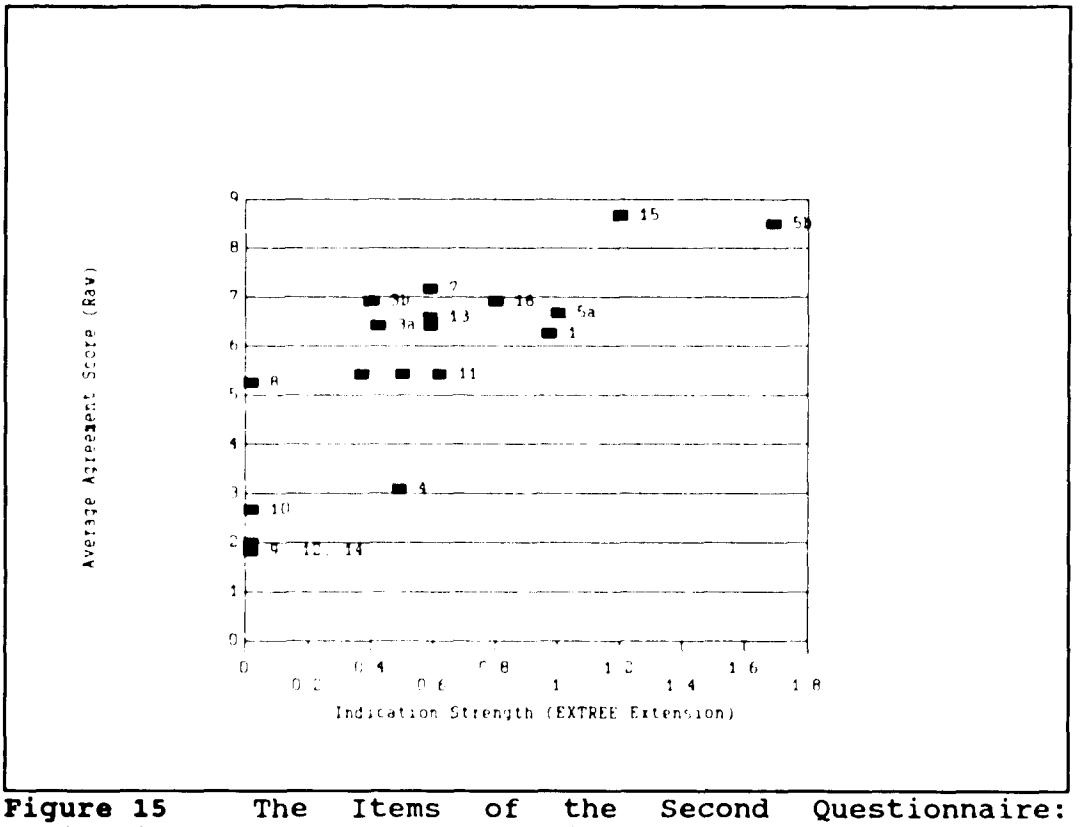

Figure 15 The Items of the Second Questionnaire: Indication Strength Plotted Against Average Agreement Score

strongly endorsed than the strength of its indication would lead one to expect. This may be due to the fact that this statement is also supported by the theory of the task structure embodied in the organization chart- that is, it is a statement that "ought to be true"- and it may be that this increases the subject's inclination to endorse the statement. It may be that similarity ratings are less influenced by this factor than is statement endorsement. In

addition, as we will see below (see page 166), it appears that this statement was substantially true at some point in the past, but was no longer true at the time of this study.

Statement #4 appears to less strongly endorsed than one would expect, given its relatively strong indication in the EXTREE output. As we will see below (page 162), these appears to be very little support for this statement, and it appears to be spurious finding.

A more illuminating perspective can be obtained by considering the responses to each specific statement in the Second Questionnaire together with the comments made in the Interviews. The following sections will treat each group of statements in turn.

Group I: (Statements #9, #10, #12, and #14)

Graph 8 is a frequency chart (giving the distribution of Second Questionnaire responses) for each of the statements in this group. It is clear that these statements were consistently and strongly rejected by the subjects.

The comments made in the interviews on statements #9, #12, and #14 were definitive:

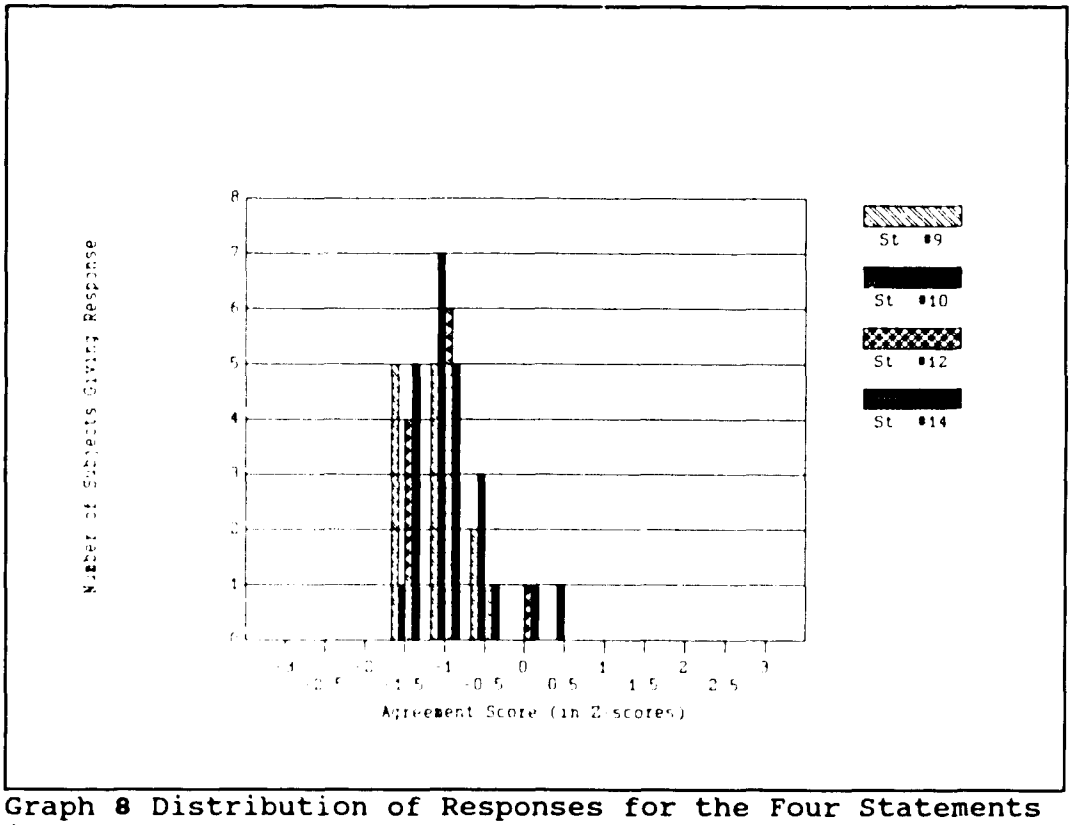

Graph 8 Distribution of Responses for the Four Statements in Group I

"I didn't see this as a working sub-unit"

"This is so false you don't have a number low enough!"

"That's too much!"

"You were just testing, weren't you?"

As a group, subjects appeared to be more willing to entertain the suggestion that there may be something correct

about statement #10, but generally concluded that it had little or no validity:

> "M_5 and M_9 aren't closely associated... did work on the education grant together."
>
> "I gave this a pretty high rating... they can talk together, whether they work together or not. [Are they a special pair?] No."
>
> "I put a two, because they weren't... it seems funny to me, they have a lot of similarities... and M_9 has been helping with M_5's research project..."
>
> "I gave this a three... didn't see it as a working sub-unit."

In these responses, it seems that the subjects are clear that these two individuals are not joined by a salient task-related group (though their work on the education project is mentioned, it does not appear to be very persuasive). Some subjects seem to explore other grounds for similarity based on personal characteristics ("they can talk together", they have a lot of similarities") but these also appear to fail to make the statement valid.

Group II: (Statements #2, #4, #8 and #11)

The individual and averaged agreement responses (in raw and within-subject z-scores) for this group of statements can be found in Table 7. As one might expect, however, the responses for this group of statements were much more mixed, and they should be treated individually.

Statement #2: "The Outpatient staff are a tightly knit group."

This statement was generally seen as being true (average raw agreement score= 6.42, average z-score= +0.30). Graph 9 shows the distribution of responses for this statement. Nine of the twelve respondents to the Second Questionnaire gave within-subject z-scores greater than 0, and their interview comments contained much agreement:

> "That's quite true... It's just a matter of contact more than anything... there is a constant back and forth, making decisions on individual patients... constant communication back and forth."

> "We communicate well in the Clinic... I have

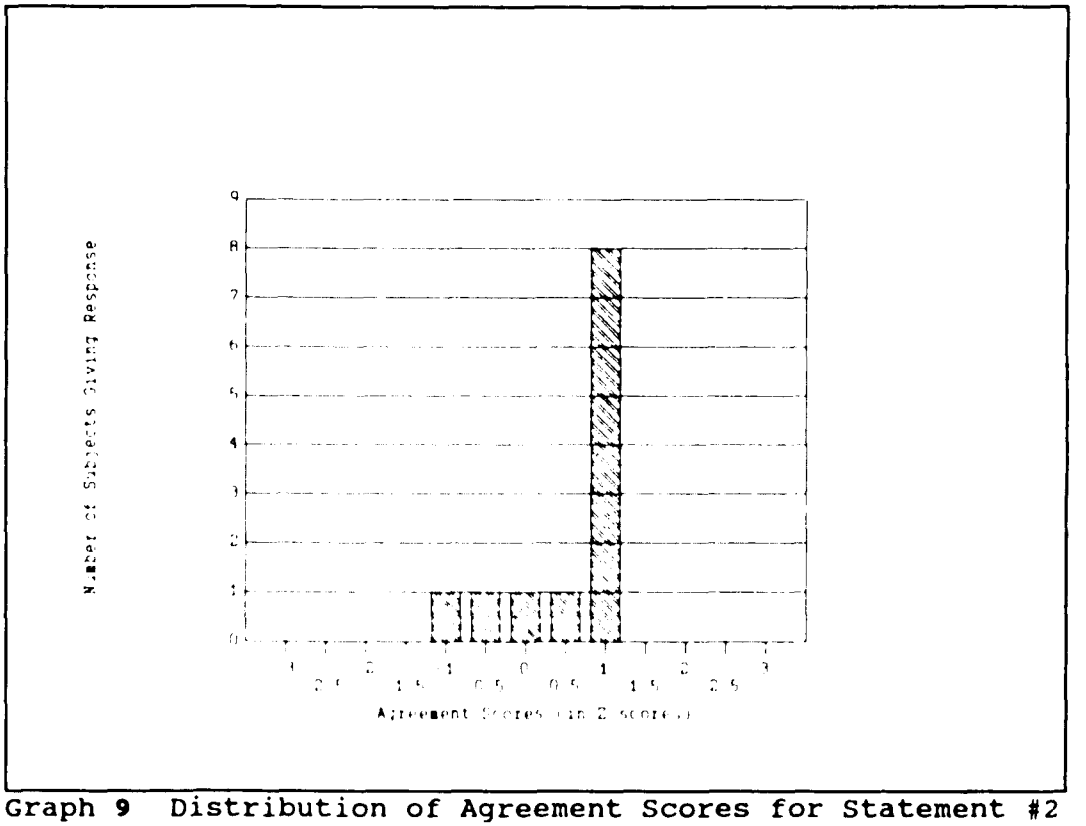

Graph 9 Distribution of Agreement Scores for Statement #2

personal relationships with everyone in the Clinic."

"We work well and closely together... physicians, physicians assistants, and the social workers."

At the same time, however, some subjects saw divisions within the Outpatient group, especially between the physician and non-physician staff:

"Not really, because the Team itself is not really tight-knit... groups within groups within groups... the docs are out a lot of the time" "This is an ambiguous question because you didn't specify what you meant by the Outpatient staff... I assume it's the people who come to the Outpatient Clinic, including the medical docs... if it doesn't, a seven is appropriate. Including the medical docs, it's two different things.""Not all of them, but I think that most try to work together. The physicians are a little more on the periphery. The Nurse Practitioner, Physicians Assistants, and Social Workers make up a more tightly knit group."

This issue will be discussed further in connection with Statements #3a and #6.

Statement #4: "The Phlebotomist is closely associated with the two Nurse Educators."

This statement was consistently rejected in both the Questionnaire responses and interviews (average raw agreement score= 3.08, average z-score= -0.89). Graph 10 shows the distribution of responses for this statement.

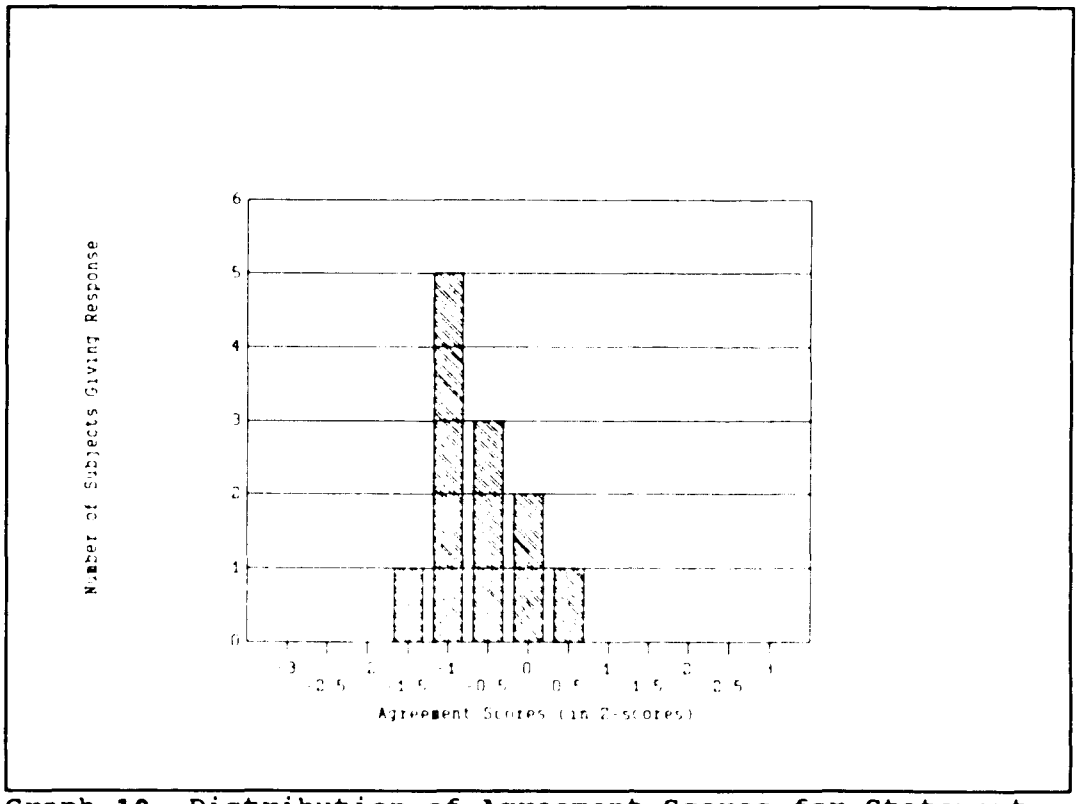

Graph **10** Distribution of Agreement Scores for Statement #4

Considerable support for the identification of the Phlebotomist [M_{18}] as an "outcast" was offered in the course of the interviews:

> "M_{18} is not very well integrated into the group... it's a difficult situation... she's fairly isolated."

> "No, I don't think they like [M_{18}]... don't think anyone does."

"I don't see [M_{18}] as associated with anybody... she's tried to be associated with them."

"M_{18} is not associated with anyone... more isolated... a self-imposed isolation."

"[M_{18}] is by herself."

Statement #8: "The Psychiatrist and the Supervising Social Worker are closely associated."

This statement evoked a mixed response (average raw agreement score= 5.25 average z-score= -0.09). Seven subjects gave (within-subject z-scores) agreement scores of greater than 0, but there were a number of dissenters, and the averaged agreement score for this statement was slightly negative. Graph 11 shows the distribution of responses for this statement.

In view of the absence of support for this statement in the similarity data, however, it is the fairly high level of agreement that seems to require explanation. In part, it may be due to the fact that this is a relationship that is called for by the organization chart (which we are taking to be "the theory of the task structure"); that is, it is something that "ought to be true." The interview data

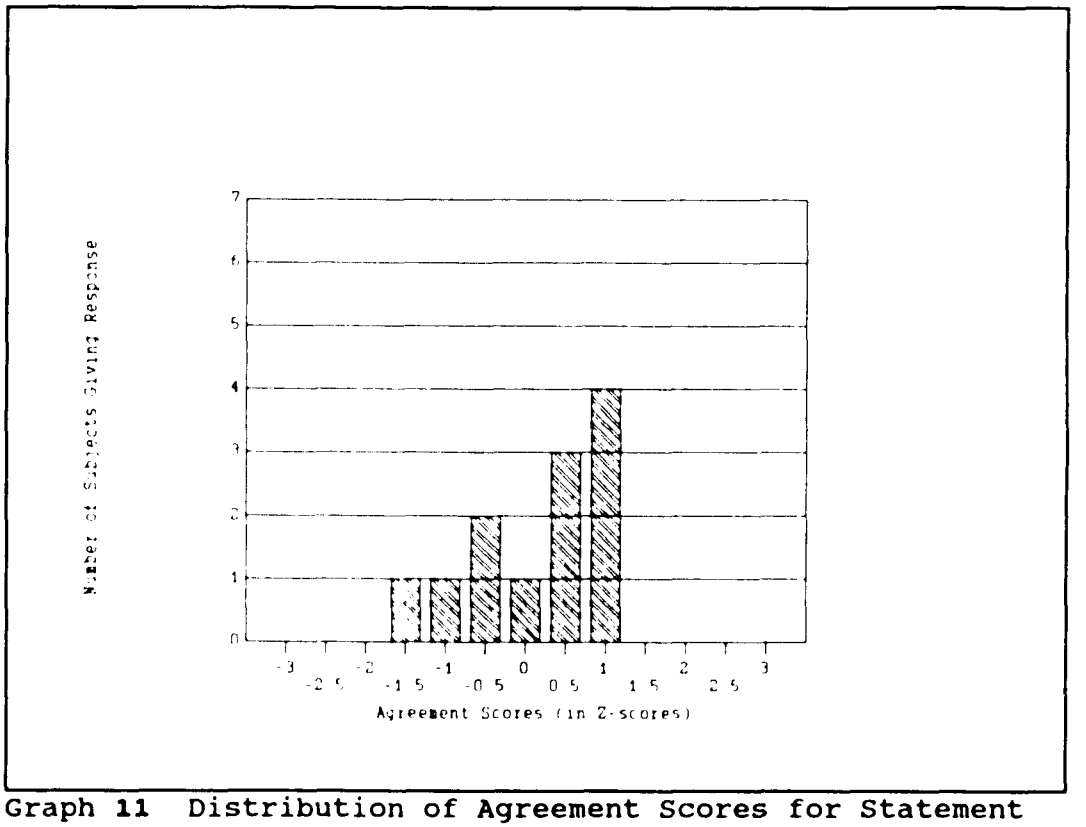

Graph 11 Distribution of Agreement Scores for Statement #8

suggests another possible explanation; that it was true at some time in the past, but is no longer true:

> "It's true, but M_{17}'s not here. If M_{17} was doing what he's supposed to be, they would be closer."

> "M_{17} and M_9... When M_{17} first came, we worked closely together, as I'm sure you know... M_{17} and I had a personal relationship... we're less close now... I found it frustrating... he wasn't here... I was angry with him."

"M_{17} and M_{9}... I gave this a seven. They tend to work closely... in the past... M_{17}, as you know is leaving... they worked closely in Clinic doing psychiatric and social service evaluations."

"I gave this a two. It's not true now."

"That used to be more true than lately. M_{17} doesn't want to do scut work... he comes to Clinic once a week... Now he's not that available."

"I gave this a seven... they were pretty close personally, even though M_{17} has pissed M_{9} off... he was a no-show for a year..."

"I put a four, toward the false end. I remember, at one time, I would have thought it pretty high... they have been pretty good friends. But with M_{17}'s not being around, everyone has been getting pissed off... M_{9} even... since he hasn't been spending time here, people's sympathies haven't been with him."

Despite the variation in agreement scores for this statement, there seems to a consensus that there had been a

close relationship between these two, but that it had deteriorated. It seems possible that subjects who agreed with the statement were referring to the past situation, while those who disagreed were referring to the present situation.

The interview data also provide very strong support for the identification of the Psychiatrist as an "outcast."

Statement #11: "The Secretary serves as a central connecting link for the Team, tying the many subunits together."

The pattern of agreement scores for this statement is similar to that for Statement #8. There is a fairly wide range of scores, and the overall average is very slightly negative (average raw agreement score= 5.42, average z-score= -0.05). Graph 12 shows the distribution of responses for this statement.

In this case, it would appear that while many agree that this is the role of the Secretary, not all agree that it is the role that the individual in question does play:

"That's definitely true. M_6 is responsible for holding the Clinics together. She does the

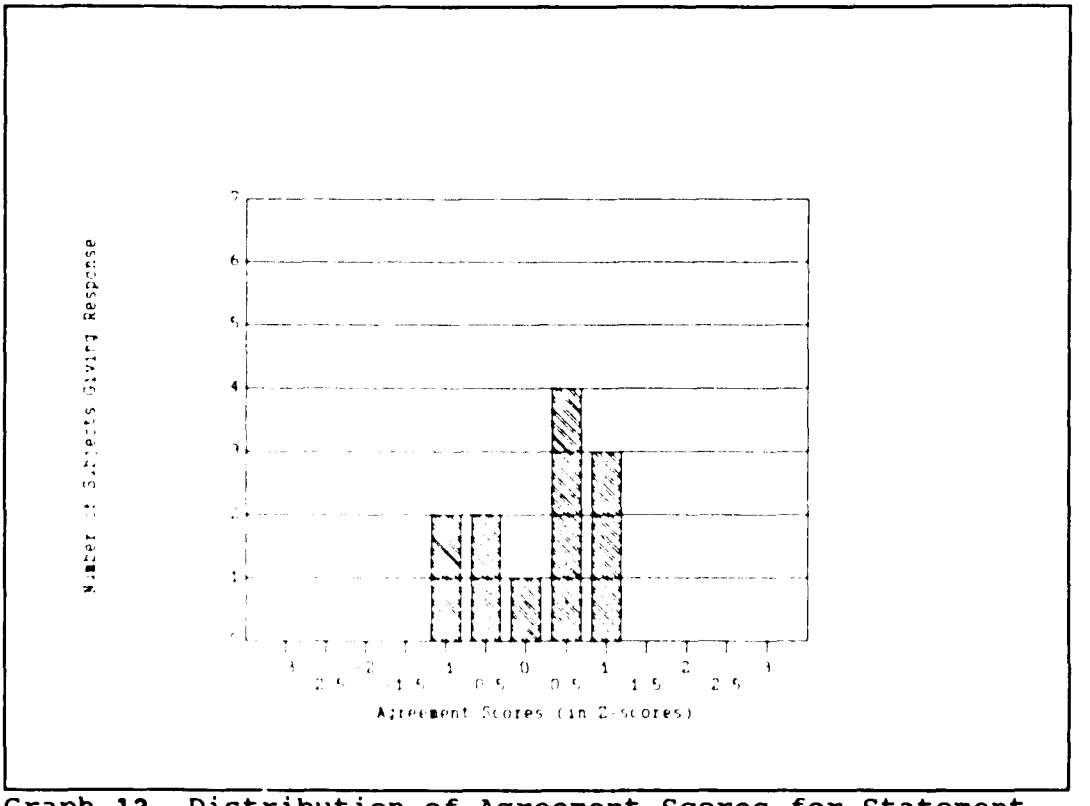

Graph 12 Distribution of Agreement Scores for Statement #11

scheduling for all of the clinicians. She negotiates with everyone on the Team. She also provides services for all of the disciplines- typing abstracts, social service reports. She works with everyone."

"I wouldn't want to overemphasize her role... by virtue of her role, she is in a central position... but not as a link which if it broke the Team would fall apart."

"M_6 came on relatively recently... I think that's the role that we envision for her... relating with all the different aspects of the Team."

"Yes, she serves that function... with some friction... it's not a good fit."

"M_6 tries, but her job is impossible. Her role is to do this, but she can't. Also, she's more connected to Outpatient than Inpatient."

"Absolutely not, although she has dealings with everyone, because she's the secretary, she does more to disrupt than anyone else."

"I see her as extremely divisive. Her role would be in the middle... that would be the role, but I don't see it occurring... she creates divisions."

"I'd put that as a one, at this point... more discord than tying together... can't say... doesn't tie them together in any way."

Group III: (Statements #1, #3a, #3b, #5a, #5b, #6, #7, #13, #15, and #16.)

These statements were consistently and fairly strongly endorsed by the subjects (see Table 7 and Graph 6).

Statement #1: "M_5 and M_1 are closely associated."

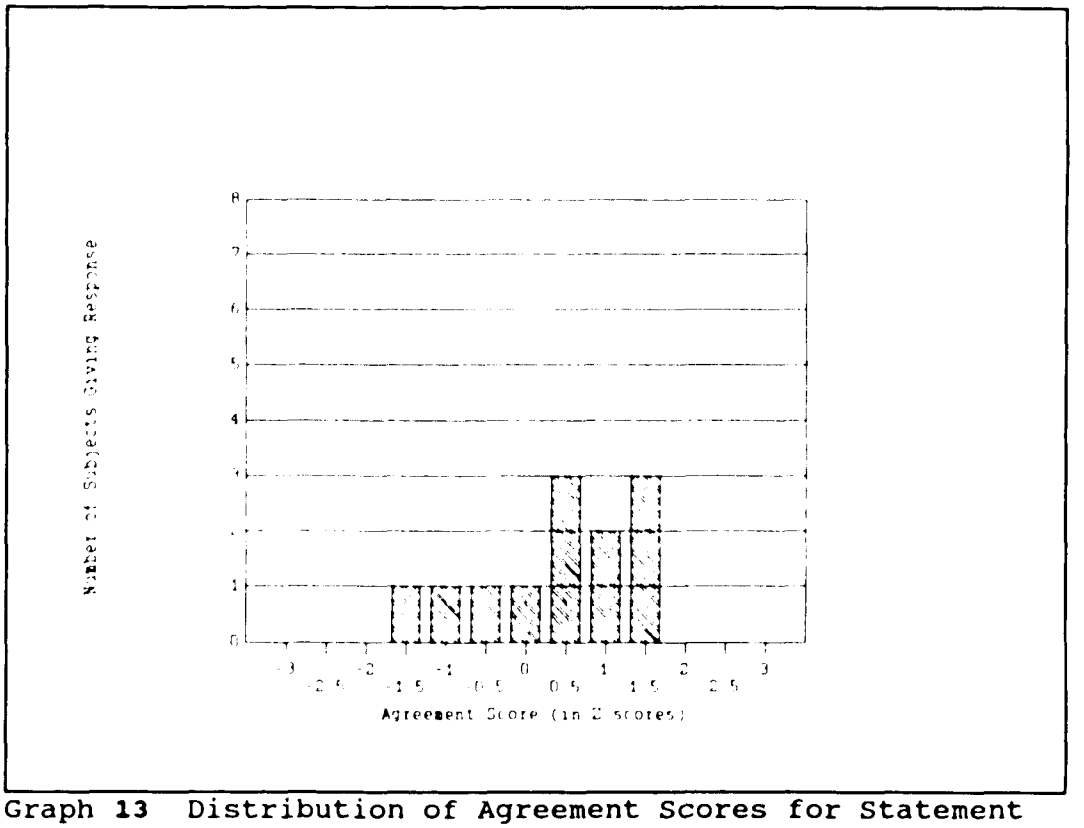

Graph 13 Distribution of Agreement Scores for Statement #1

Though the average agreement scores were only slightly positive (average raw score= 6.25, average z-score= +0.19), most of the subjects agreed with it. Graph 13 shows the distribution of responses for this statement. Eight of the twelve respondents gave agreement scores greater than 0, though there were three dissenters.

The interview subjects (who, unfortunately, did not include any of the subjects who disagreed with this statement) generally confirmed the notion of a strong personal connection between these two, and suggested that this was not a task-related pairing:

> "I rated this as highly true. They are very close, almost brotherly as M_1 would put it... in terms of role, much, much less. (What is M_5's role?) Not clear... he was brought in to get inpatient up and running... he's overcommitted."

> "I see this as... the two have been around for a long time... it's more of a personal relationship rather than a structural one... one that is built into the structure of the Team...The fact that M_5 is on the Team has more to do with the fact that he and M_1 have been kicking around here for ten years... it's a personal relationship that has evolved..."

The interviews also provided unequivocal support for the identification of M_5 as an "outcast":

> "I don't know anything about M_5. (What is his

role on the Team?) I don't know. He's a part-time physician... I don't know what he does."

"Not clear... he was brought in to get inpatient up and running... he's overcommitted."

"He's on the team 1/2 time as an attending... really only comes to one clinic. I've gotten to know him better through the education work that we've done together. That's not a Team activity."

"M_5 is also functioning as a half-time attending... he's theoretically providing attending-level support in Clinic..."

"On the Team, he comes to Clinic one day a week... he was supposed to have a more active role as a facilitator on Inpatient, but it has not materialized."

"I don't really know why he [M_5] is on the X Team... he goes to Clinic only very occasionally. In order for him to be on the Team and do as little as he does, there must be some sort of longstanding friendship between the two..."

"M_5's role is very minimal, in terms of the Team... M_1 brought him on with promises that he would take a role on the interactions between the Team members... and work on the Inpatient side. Now all he does is one Clinic a week."

All in all, the picture of this person presented by the similarity data, of an individual with a reasonably strong, but not task-related, connection to the Medical Director but no other connections to the Team, appears to be true to life.

Statement #3a: "The Outpatient staff contains two distinct sub-groups... [one of which is] the three attending physicians."

This statement was endorsed by seven of the respondents, but its averaged agreement scores (average raw score= 6.25, average z-score= +0.22) were relatively small. Graph 14 shows the distribution of responses for this statement. Its companion statement (#3b) was agreed with to a slightly greater degree and, in fact, Statement #2, which is at least implicitly contradictory, also received a higher agreement score.

The interview data suggests that two different currents

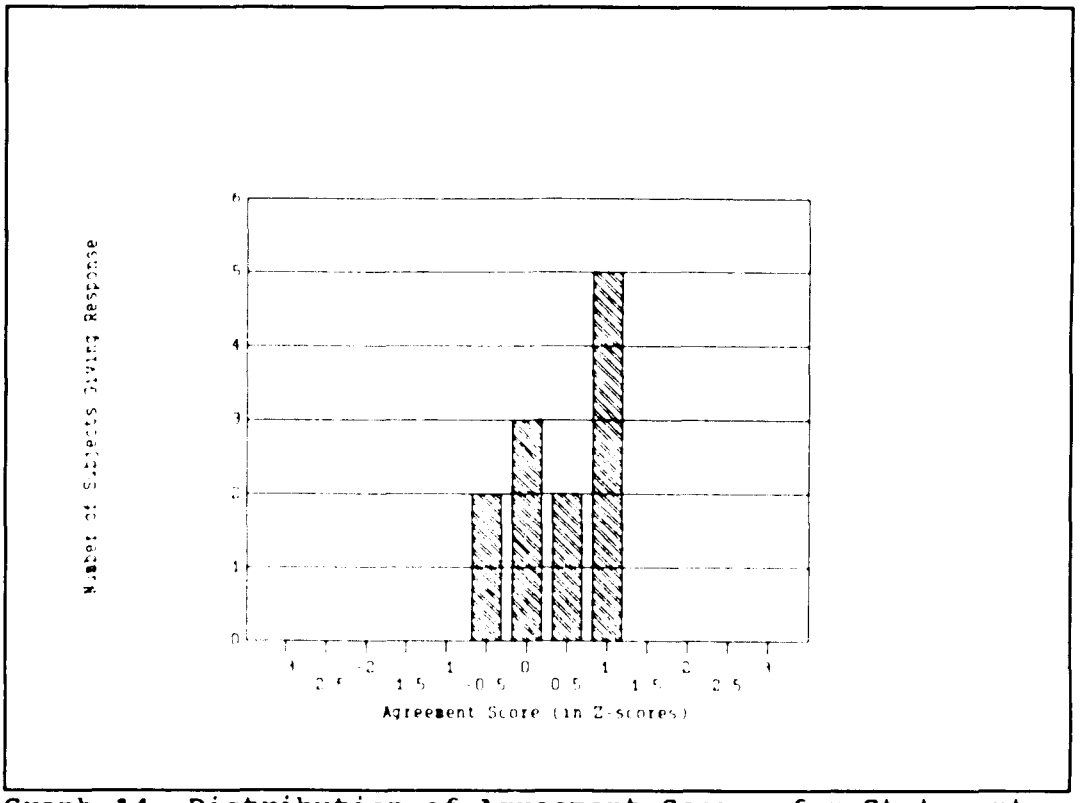

Graph **14** Distribution of Agreement Scores for Statement #3a

may be at work here. On the one hand, the physicians do seem to appear to be set off from the non-physicians in the group, because of their discipline and because of the conditions under which they work. As we will see in relation to Statement #6 below, there is little disagreement that the physicians are a distinct subgroup within the Team. On the other hand, it appears that one of the individuals in the proposed sub-unit is actually somewhat different from the other two, and this seems to have detracted from the plausibility of the statement:

"I don't know if I'd say this... I never see M_3... M_2 and I have Monday clinics together, plus we are working on a couple of research projects together."

"I don't see M_4 as part of that group... M_2 and M_3 maybe... not M_4... M_2 and M_3 yes... they're the frontline physicians... they're more a part of the Team than the administrative types."

"I gave this [**3a**] an eight... we tend to function in a similar vein..."

"M_3, M_2, and M_4 are sort of bound by their professional role... I might downgrade that... they don't think and act like a group."

"Kinda... M_4's apart from the rest of them... but I had to say an eight... M_2 and M_3 share an office... Don't know what they [the physicians] do... they say they work 13 hour days... I don't disbelieve them but... I think M_1 must have told them, at some point, that all they had to do was go to a couple of Clinics...M_4... I don't know what he does. We call M_4 the "Prince of the Team," M_1 is the "King."

"M_2 and M_3 are very involved. M_4 is not. All three strike me as being very cool temperaments."

"The docs are some kind of group... M_4's not a part... M_2 and M_3, yes."

"They are all physicians, that is the tie that binds them all together. M_4 is a little bit separate- that's because he's above them, as Assistant Director."

When the similarity data are reexamined, in light of these comments, it seems that Statement #3a was poorly phrased. The similarity data place M_4 in a somewhat different place and do not necessarily support the existence of the trio suggested by this item.

Statement #3b: The Outpatient staff contains two distinct subgroups... [one of which is] the Social Worker and one of the Physicians Assistants."

The statement was fairly consistently and fairly strongly endorsed by the respondents (average raw agreement score= 6.92, z-score= +0.47). Graph 15 shows the

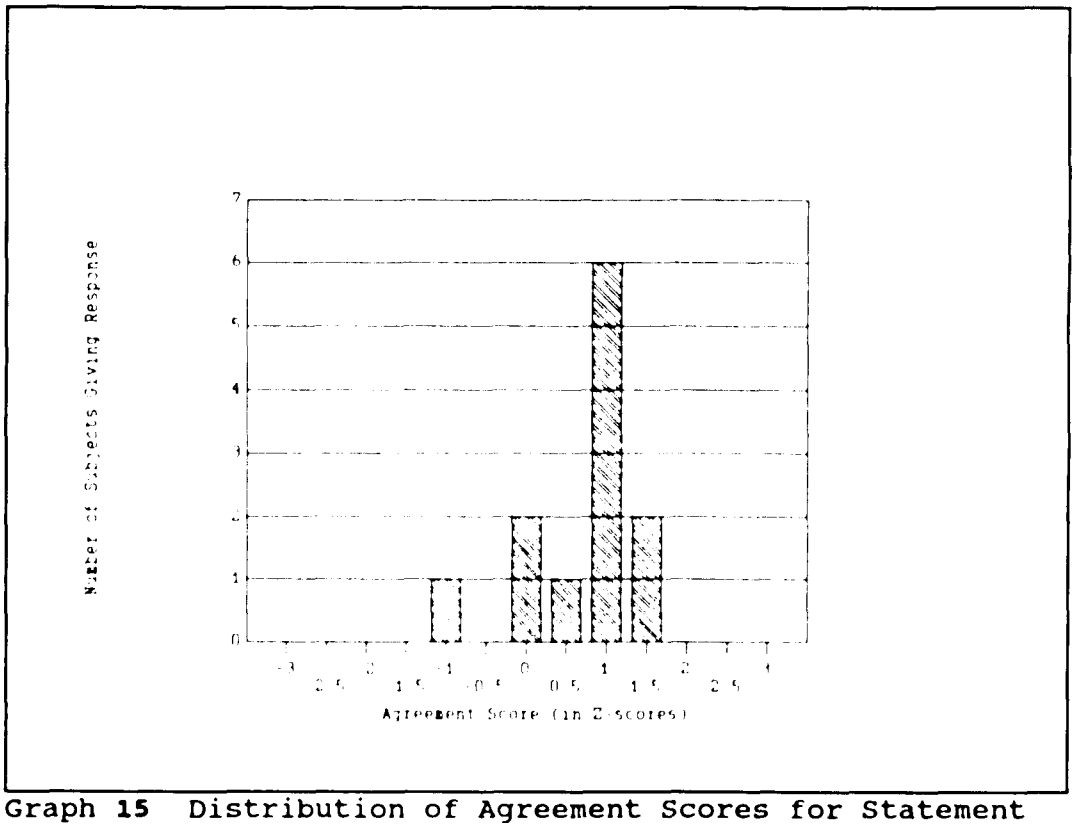

Graph **15** Distribution of Agreement Scores for Statement #3b

distribution of responses for this statement. Only one strong dissent was expressed.

In interviews, most subjects were clear that this was a recognizable special pairing within the group:

"M_8 and M_9 do a lot of work closely together."

"I gave this a six... We're friends... wouldn't say we're a subgroup... don't see us as being cliquish... (What about a "special relationship"

as versus "cliquish?") Yeah, that could be said."

"The social workers... they're very tight... I said a nine... they're good personal friends."

"Yeah, they're close friends. M_9 and M_8 have an emotional way of relating."

"Absolutely true... M_9 and I are very close. We've been here the longest. I've been here for 2 and 1/2 years... M_9 was here first... about 3 years. We did the interviewing and hiring for the Team members. M_{10}- a nurse practitioner- was hired to work on the Team, but only temporarily. She is now across the street. It was known from the beginning that she would head the X Research Project. She works in the Clinic a few days a week, but not much else. We get along well professionally, but we aren't as close as M_9 and I are."

The dissenting subject acknowledged the closeness between these two, but did not see them as being a distinct subgroup:

"The substructure that I see is M_8, M_9, M_7, and I-

we're friends. M_8 and M_9 alone don't seem that distinct."

Though many of the comments tend to describe the relationship in personal (rather than role) terms, there are also suggestions that these two individuals have a tacit leadership role for the Outpatient section of the Team. As one person put it:

"They provide... I don't want to say clerical... but they make things run, see that the paper flows, that labs get checked... they work in a similar capacity- M_9 in Social Work, M_8 in clinical care. I would tend to include... it's interesting... there is another physician's assistant and the nurse practitioner... but it's true that M_8 and M_9 are functioning on a different level... not exactly supervisory... they tend to take more responsibility than other members of the Team."

Statement #5a: "The Social Work/Psychiatric staff group contains two distinct subgroups... [one which is] the Supervising Social Worker and the Pediatric Social Worker."

The statement, which links two staff members who belong to the same discipline but work in distinctly different

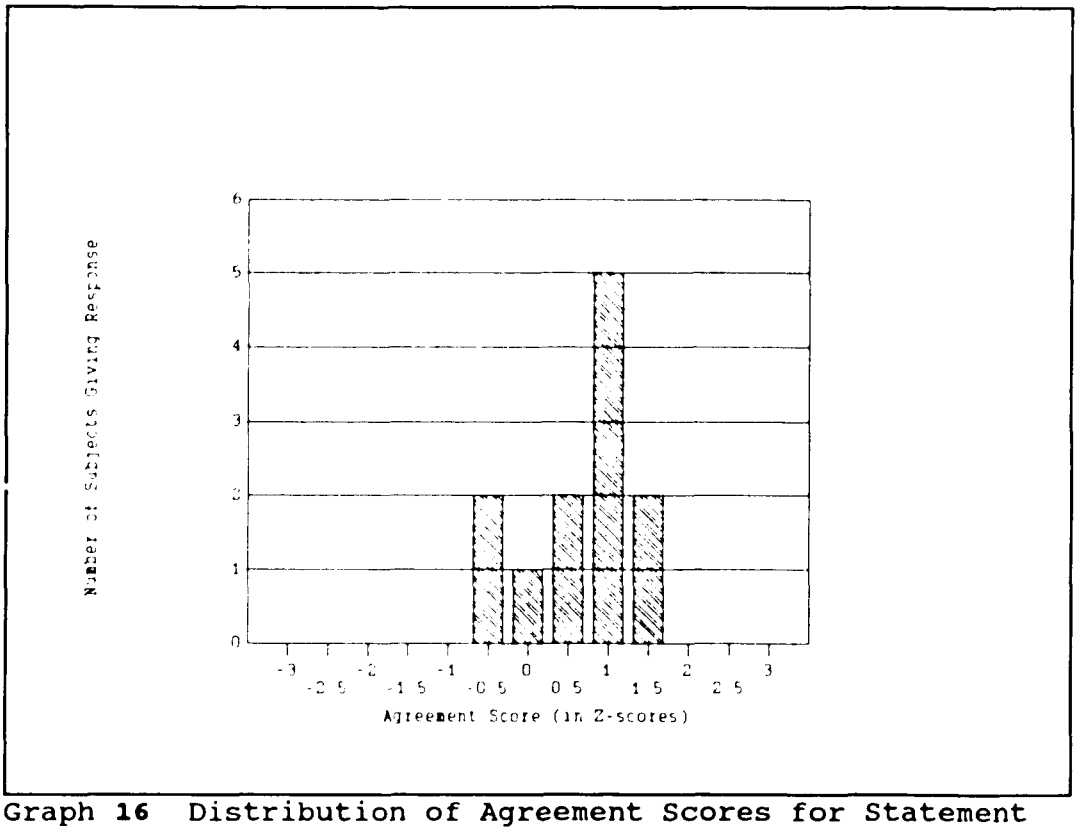

Graph 16 Distribution of Agreement Scores for Statement #5a

settings (Adult Outpatient and Pediatrics), was given moderately high and consistent agreement ratings (average raw agreement score= 6.67, z-score= +0.45). Graph 16 shows the distribution of responses for this statement.

One of the two subjects who gave it agreement scores below 0 was interviewed, and indicated that though he had not seen this as a relevant task grouping, he did acknowledge the close personal connection:

> "Why these two? Do they work together... they don't... M_{11} is on her own, M_9 too... (They're not a pair?) In terms of work, no... in terms of friendship, yes... they are very close friends. M_9 is the only one M_1 is really close with on the Team."

The other subjects also acknowledged the strong personal relationship:

> "With M_9 and M_{11}, I notice a camaraderie. They don't work as closely together, but there is a definite sense of unity, a strong sense of connection."

> "M_{11} and I are close friends, we socialize outside

of work... we know each other more than anyone else... we work on projects together- presentations and so forth... we don't actually work together."

"I put a six here... I don't like the idea of these "subgroups"... M_9 is more of a friend of mine than the others on the Team... I don't like the idea of excluding people... ("Subgroups" don't have to be exclusive- it just means that some people have especially close relationships.) Oh, then I'd make it higher, probably an 8."

A few saw a task relationship as well:

"M_{11} and M_9 are very close- no question about that- they also work very closely together."

Statement #5b: "The Social Work/Psychiatric staff group contains two distinct subgroups... [one which is] the Inpatient Social Worker and the Inpatient Case Worker."

This part of Statement #5 received emphatic and unanimous agreement (average raw agreement score= 8.50, z-score= +1.08). No subject gave it an agreement score below

.80. Graph 17 shows the distribution of responses for this statement.

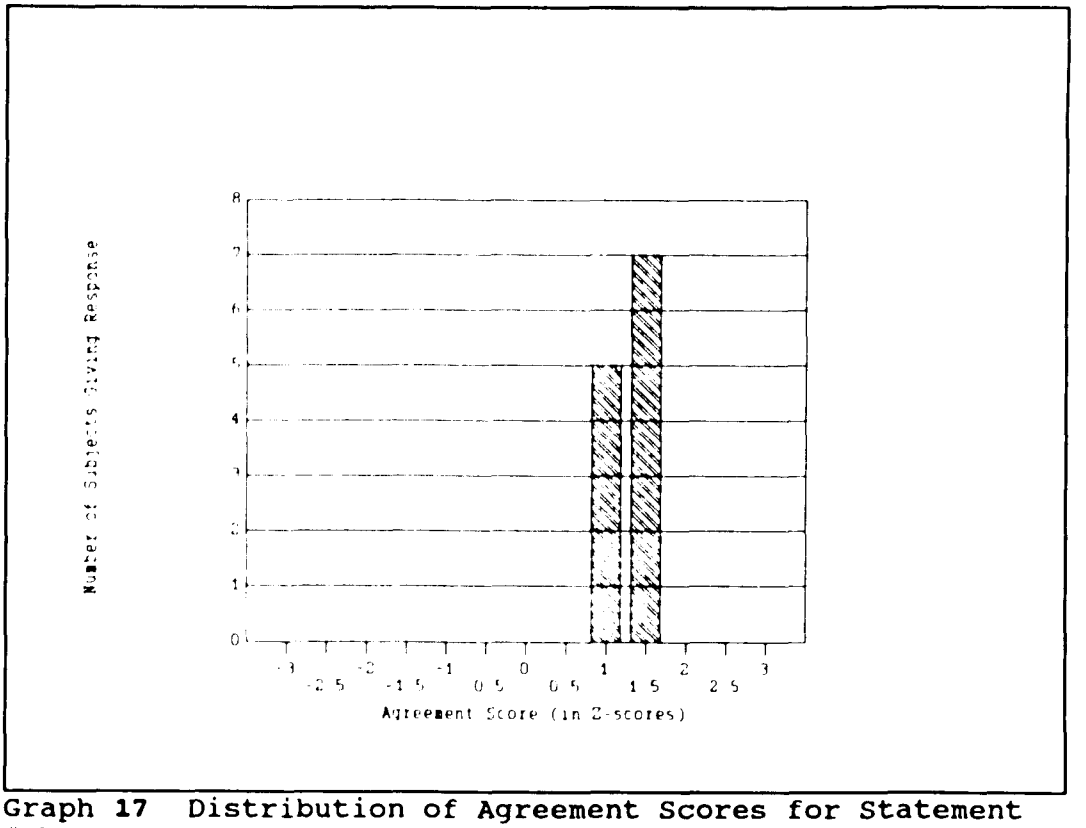

Graph 17 Distribution of Agreement Scores for Statement #5b

In this pairing, it seems that the boundaries of a close personal relationship and a close working relationship coincide:

> "M_{14} and M_{13} always do everything together... work on clients, have lunch together."

"Absolutely... they don't go anywhere without each other, hardly."

Overall, it would appear that the basic idea that underlies the two parts of Statement #5, that there are two distinct subgroups within the Social Work staff, seems to be valid.

Statement #6: "In general, the physicians on the Team have more in common with each other than they do with the non-physician staff of the units (e.g. Inpatient, Outpatient) to which they are assigned."

This statement, the first of two that attempt to touch upon some of the conflicts that are intrinsic to interdisciplinary work groups, received strong agreement from the subjects (average raw agreement score= 7.75, z-score= +0.78). Graph 18 shows the distribution of responses for this statement. Only two subjects gave agreement scores less than 0, and seven gave agreement scores greater than 1.0.

In the interviews, it seemed that most subjects were strongly influenced by the salience of the physician role, in spite of the recognition that the physicians were, as

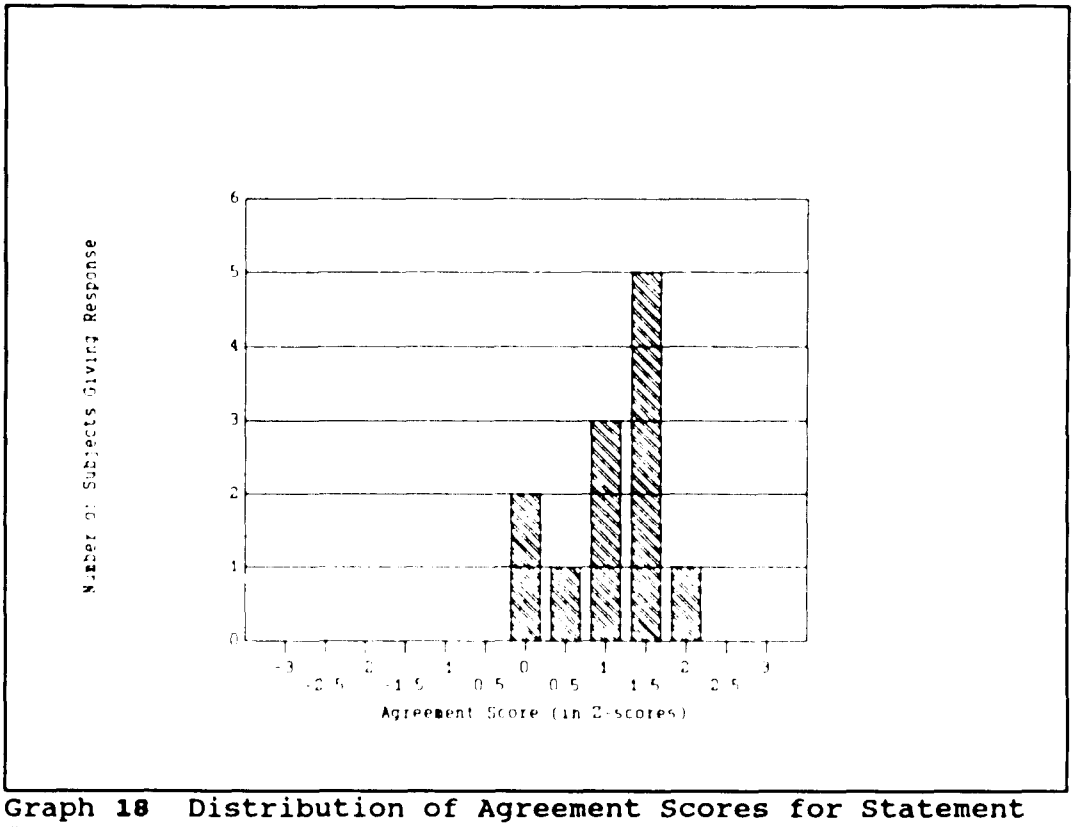

Graph **18** Distribution of Agreement Scores for Statement #6

people, a diverse group:

> "I had a hard time separating personal from professional... they have the physician role in common... but not personalities..."
>
> "We functionally tend to operate in the same mode... it just seems intuitively true."
>
> "I don't see or hear of them relating to each other, but they have a "generic doctorhood" with a

capital 'D'."

Some subjects attribute the distinctiveness of the physician group to differences in the ethnic or gender composition of the group:

> "The non-physicians are a diverse group. To be frank, the physicians are white, middle-class, and male. The non-physicians are a variety."

Most often, however, subjects cited differences in responsibilities and privileges that accrue to the physician role, especially in an academic health center:

> "The docs do Clinic, but that doesn't last very long... technically it's supposed to last till three, but it goes to four... We don't have a doc in this office... they stop in occasionally..."

> "Yeah... we once had a big blowup [in the support group]... about the physicians being separate from the Team... a long meeting ensued... then M_5 stopped by and shed some light on it... he explained the dual appointments, said that the doc's were hired on tenure track lines, have to publish... and we said Oh!..."

Statement #7: "These two subunits of the Social Work/Psychiatric staff are each more closely associated with their places of work (Inpatient, Outpatient) than they are with their disciplinary base (Social Work/Psychiatric)."

All subjects gave positive responses to this item on

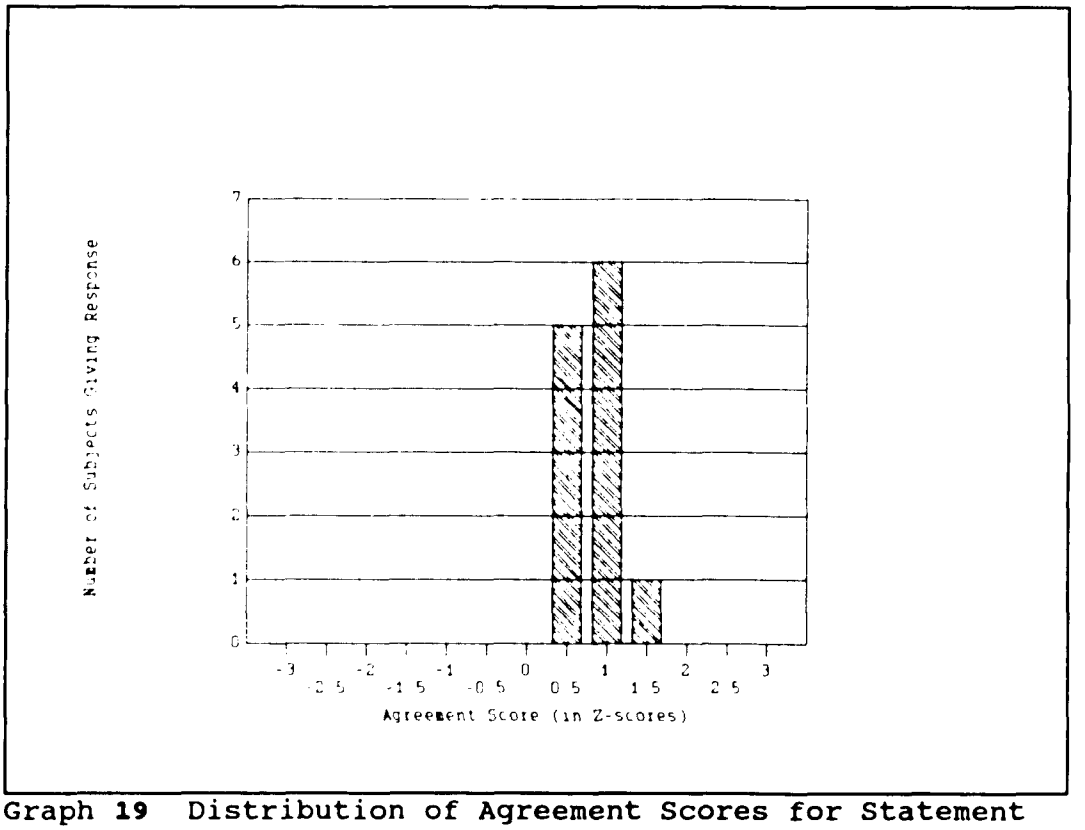

Graph 19 Distribution of Agreement Scores for Statement #7

the Second Questionnaire (average raw agreement score= 7.17, z-score= +0.58) and, as Graph 19 indicates, many subjects were strongly positive. In the interviews, however, many subjects indicated that they had not been sure what the statement meant. In several instances, they seemed to think that it referred to the affiliation with the Hospital's Department of Social Work, and seemed to be unaware that there was a Social Work/Psychiatric unit within the X Team:

> "I think that we're distinct... more closely associated with the X Team than with Social Work. M_{11} and M_9 are very much X Team workers, so are the others."

> "What does this mean... I have no idea. It's not really a real unit. One person, a psychiatrist, resigned. Prior to resigning, he was non-existent anyway. That division is false."

This sort of response can be interpreted as a confirmation of the basic idea that underlay this statement- that is, that the Social Work/Psychiatric group described by the Medical Director is not a meaningful subgroup ("is not psychologically real") for the Team. Furthermore, some subjects spoke very directly to the issue (whether they understood the question or not):

"It is absolutely true, but I gave it a seven when I marked it... what was I thinking of? I try to know what's going on in Inpatient, but I don't, really... I would have marked it higher... "with their disciplinary base"... we don't really have... we're not bound together by being members of a common discipline."

Statement #13: "The Supervising Social Worker works closely with the physician leadership of the Team."

This statement received a moderately high average agreement score (average raw agreement score= 6.58, z-score= +0.42) and, as Graph 20 shows, the agreement was relatively consistent.

The comments in the interviews suggested that the Supervising Social Worker has, perhaps because of the absence of the Psychiatrist and the part-time commitments of the physicians who are the nominal leaders of the Team, served as a de facto leader for the non-physician staff who

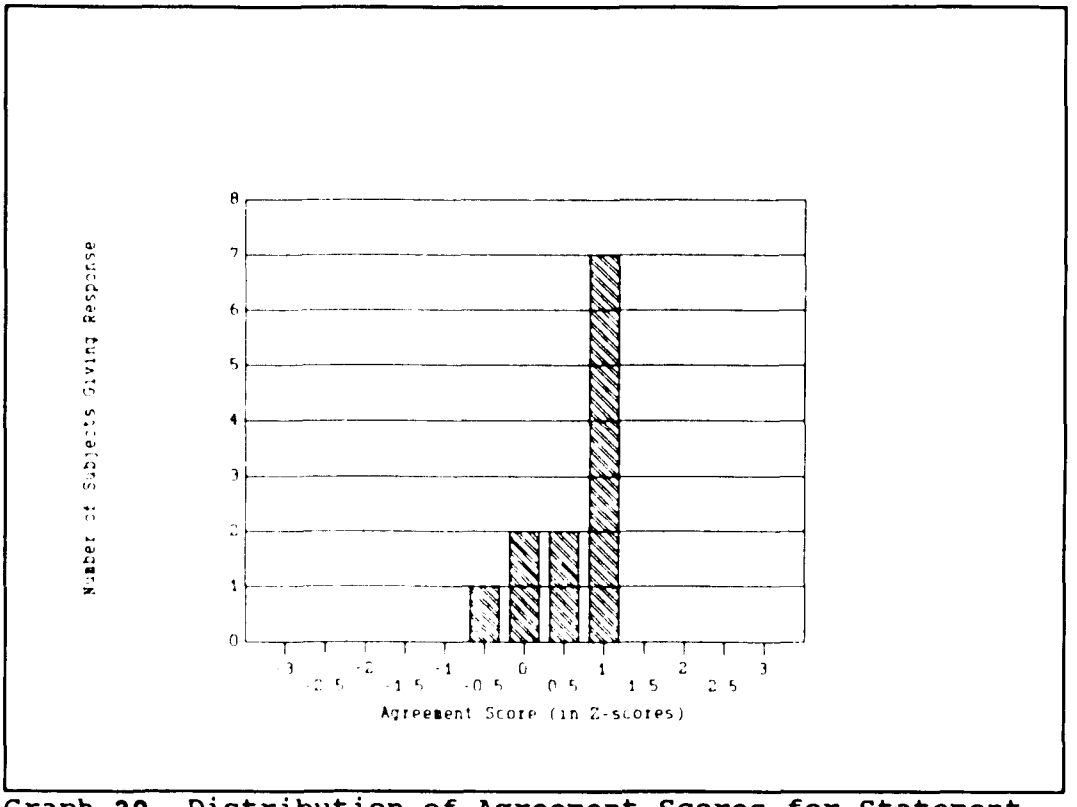

Graph **20** Distribution of Agreement Scores for Statement #13

work in the Outpatient Clinics[20]:

> "Yes... she is responsible, at least for Outpatient... she's always in Clinic which is where the primary professional activity occurs,and she works very easily with physicians."
>
> "Yes, she has emerged as a leader."

---

[20] It is noteworthy that, after the completion of this study, the Supervising Social Worker has become the administrator of the Team and left her clinical service role.

"Yeah she does. As a social worker, she's very good at working with medical information, more than other social workers. She has a whole supervisory role as a supervising social worker... takes on administrative tasks... able to see the issues of the whole X Team more clearly..."

Even the subjects who gave the statement low agreement scores appeared to see some validity to it:

"I gave this a two, but... if anyone is the link between social work and the staff... and the rest of the Team... it's M_9... and the physicians are part of the rest of the Team."

Statement #15: "The Two Nurse Educators form a distinct sub-unit within the Team."

The statement received the highest average agreement scores (average raw score= 8.67, average z-score= +1.14) of all the statements. As Graph 21 shows, the agreement was unanimous as well as strong. Interview comments were equally clear, and spoke of the ways in which this pairing had a distinct task and operated in a distinct setting:

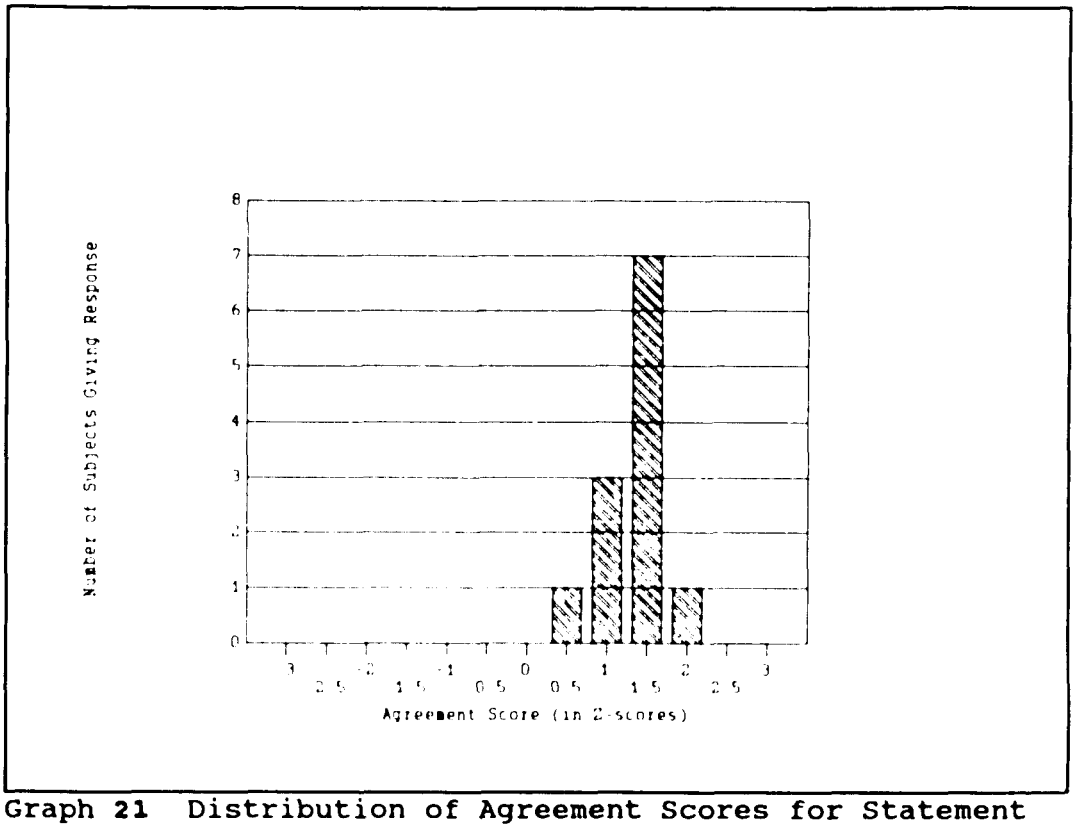

Graph 21 Distribution of Agreement Scores for Statement #15

"Absolutely. They're separate from us... a separate office. They work well with the rest of us, though."

"They function in an area that no one else does. They are more separate, physically and functionally."

"They're unique, both because of location, discipline. More so since they moved to a new building... now they are God-knows-where in A

Building."

"Absolutely. You almost never mention one without mentioning the other. Long ago they seemed very separate [from the rest of the Team] now they seem much closer in the thing, which is nice."

Statement #16: "One of the Physician Assistants and the Coordinating Manager form a distinct sub-unit within the Team."

This statement received fairly high agreement scores (average raw score= 6.92, average z-score= +0.50). As Graph 22 shows, most subjects agreed with it, though there are two dissents. In the interviews, a consistent account emerged. The two individuals had very different tasks in two different sectors of the Team, but were drawn to each other and developed a close personal relationship. In time, apparently as an outgrowth of this personal relationship, they began to take on and share some common tasks:

"The personal relationship seemed to be stronger than with other members of the Team. On the hierarchical schematic they are not really related, but personally and functionally they volunteered... they work in totally different

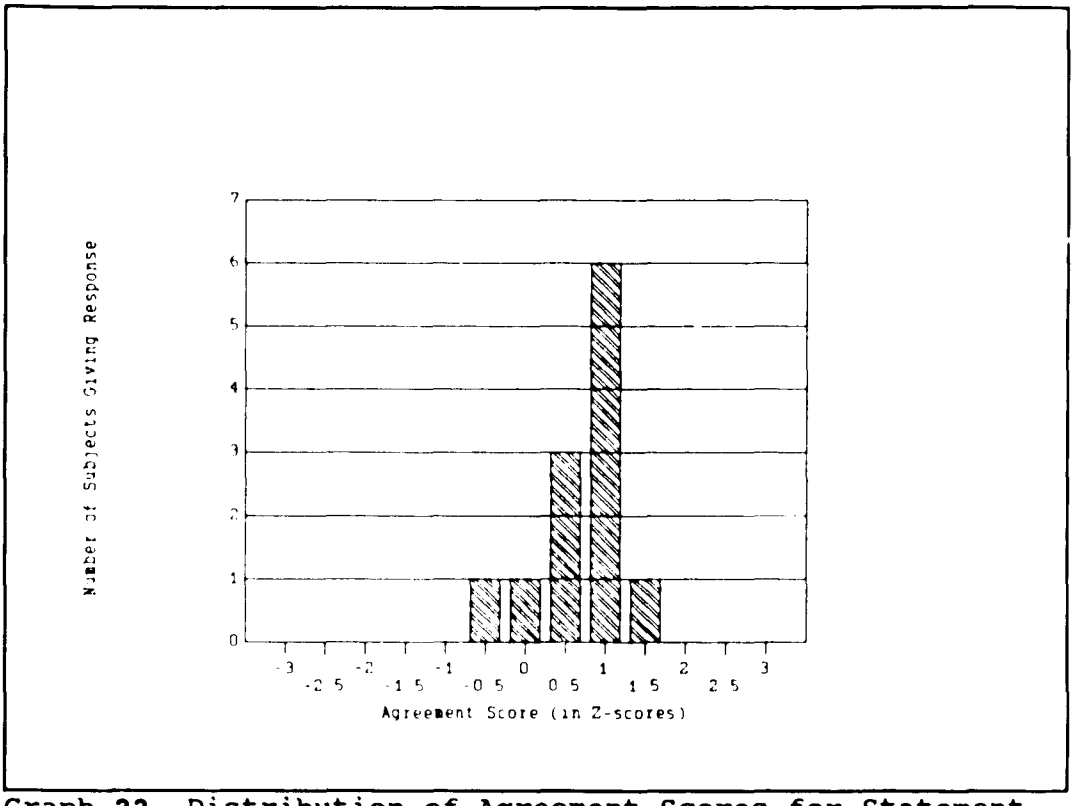

Graph 22 Distribution of Agreement Scores for Statement #16

places... the work evolved out of their own interests..."

"They're a subgroup on a personal level... M_7 helped M_{12} do some stuff."

"Yeah, more in terms of personality than work... they just gravitated toward each other."

As for the two dissents, it appeared that one was not aware of the close personal relationship, and both tended to

emphasize the lack of a formal work relationship:

> "They're close personally, but not in terms of work. M_7 helps with some of the counseling work that M_{12} does."

> "M_7 and M_{12}- I wouldn't have put them together... I don't know why... they do do blood testing on the Inpatient units."

Several subjects mentioned the fact that these individuals were the only identified homosexuals on the Team:

> "M_7 and M_{12} are both gay, and they do have a lot of common issues around that..."

> "They were good friends... the only two gay men on the Team... the only two out there gay men..."

V. DISCUSSION

In this discussion, I will take up three topics. First, I will attempt to characterize the findings of the two studies reported here in a general way. The thrust of this section will be that the method proposed here does, in fact, provide a way of determining the actual distribution of sentience within an organization at a given time. The second section will be devoted to a detailed description of the method and of the many practical problems that still remain to be resolved in its application. In a sense, this second section can be seen as a guide to the current "state of the art" in the application of the method. Finally, I will turn to the question of what, if anything, the results reported here can tell us about the nature of the representations that individuals use to represent their work environments.

The Findings

The two studies reported here were conducted in order to find out if it was possible to determine the actual distribution of sentience within an organization by: 1) asking the members of the organization to rate the similarity of members to each other; and 2) applying the

currently available algorithms to the resulting data to construct configurations which can account for the similarity data. These configurations, it was thought, might reflect the distribution of sentience.

The logic that underlies this method was based on the analogy to the relationship between a road map and its associated mileage chart, which suggested a method for (partially) reconstructing the original map from a set of interpoint distances. In this method, one fits the set of distances to several different types of map, in an attempt to find a type of map that permits a reasonably good reconstruction of the distances. Once such a map is found, one can tentatively conclude that it is an approximation of the original map. In the case of the road map, we may assume, from the outset, that we know what type of map it is, and that we know what the distances on the mileage chart represent. In such a case, all that we need to reconstruct is the exact locations (on the map) of the objects whose mutual distances have been given in the mileage chart. The case that we have been considering in these studies is different in that we have not been able to assume that we knew what the "distances" that we found in the similarity ratings represented, or that we knew the kind of "map" from which they were derived. In this respect, our task is more difficult in that we not only have to reconstruct a map

which provides an acceptable account of the similarity data, we have to provide an interpretation of the map as well.

The question then becomes, do the "maps" that we have constructed from the similarity data gathered in the studies reported here depict the sorts of phenomena that Miller and Rice called "sentience" or do they, instead, depict some other sorts of facts about the organizations and individuals from which they were drawn? On the whole, the answer seem to be that the configurations that were constructed from the similarity data do provide a very "true to life" portrait of the actual distribution of sentience within the two organizations studied, at the time that the data was gathered.

In general, there are two types of evidence that support this contention. The first is the strong general resemblance, in both of the studies reported here, between the configurations derived from the similarity data and the organizational charts for the groups from which the similarity data was drawn. The best illustration of this correspondence is the location of the overall leader of the organization. In the MDS analyses of the similarity data from both organizations, the leader was located in a position that was at one extreme of one of the dimensions (interpreted as a "status" dimension) and at the center of

the plane defined by two other dimensions. This location, above everyone but equidistant from the major subdivisions in the organization, corresponds very closely to the location defined for the leader in the organizational chart. Furthermore, many of the major elements of the structure that appeared in the output of the computer algorithms correspond to the major divisions in the organization. For example, in the Pilot Study, three of the major functional divisions (General Administration, Service, and Educational) appeared consistently, in all of the configurations, as distinct groupings.

The organizational chart describes the groupings of roles and tasks that have been defined as part of the division of labor. "Sentience" refers to the groups with which the individual members have identified themselves (and, presumably, are identified by others). We would expect that, in order for the organization to function at all, there must be some degree of correspondence between the two and it seems reasonable, therefore, to view the resemblance that we find between the two in the organizations studied here as indicating that the maps that we derive from the similarity data are concerned with representing the same sorts of facts that the organizational charts are.

At the same time, however, there are many respects in which the organizational charts and the configurations derived from the similarity data differ. Though it may seem paradoxical at first, I believe that these discrepancies also constitute persuasive evidence that the reconstructions from similarity data reflect the true distribution of sentience within the organization. We must remember that the organizational chart is a only theory about the way in which the group is organized, and that it may not correspond to reality in every respect. Both of the studies reported here produced instances where the picture derived from the similarity ratings appears to be more "true to life" than the picture derived from the organizational chart.

Probably the most striking, and consistent, such finding is the identification of the "outcasts." These individuals, whose location in the similarity-derived configurations was very different from that which would be expected on the basis of their formal role in the organization, were individuals who were widely viewed (as confirmed by the subjects interviewed in the Second Study) has having difficulty taking up their roles or in fitting in to their assigned niche in the organization. In other words, the discrepant location in the output configuration seems to reflect a real aspect of their situation in the organization, an aspect which was not reflected in the

organizational chart.

Another such finding is the way in which long-standing professional/personal relationships between individuals are reflected in the details of the tree structures. In both studies, but most strikingly in the Pilot Study, the structure of some of the subclusters (e.g. the General Administrative and CEO subclusters) reflected the length of association between the leader of the subcluster and the other members. Again, the configuration reveals a real (and meaningful) fact about the organization which is not made manifest by the organizational chart. Furthermore, it is a fact that seems quite compatible with the general notion of sentience- it suggests that individuals who have worked closely together for a long time become identified with each other, in the minds of the members of the organization.

A third finding is the discovery, in the Second Study, of a task group that was not included in the organizational chart, but which appears to reflect, at least in part, a real sub-task. In the Second Study, the Coordinating Manager and one of the Physician's Assistants emerged as a rather salient pairing (in the reconstructions of the data from the First Questionnaire and the confirmations of the hypotheses in the Second Questionnaire), though they occupied quite disparate locations in the organizational

chart. In the interviews, two distinct, but not inconsistent, explanations were suggested for this. First, it was widely acknowledged that the two individuals had an exceptionally close personal relationship. Second, it appeared that they had, apparently on their own initiative, taken on a common task which was not explicitly acknowledged in the organizational chart.

A fourth such finding is the non-reality of the Social Work/Psychiatric subunit in the Second Study. In this instance, the organizational chart drawn by the Medical Director included a major functional sub-unit that did not appear to have any impact on the similarity ratings of the subjects (in other words, the proposed cluster had no "psychological reality" in the minds of the subjects). Though this discrepancy appeared troublesome, at first, the data from both the Second Questionnaire and the Interviews indicated that the proposed grouping did not appear to play a significant role in either the working life or psychological life of the group. In other words, the conflict between the organizational chart and the similarity data was resolved in favor of the similarity data.

There are many other examples of such discrepancies between the theory of the organizational structure and the actual structure, as revealed by the similarity data. The

existence of such discrepancies, and the fact that the similarity-derived representations are so often correct, suggests that the technique might eventually have some utility as a method of organizational diagnosis. The following section will discuss the current state of the art in the use of the technique and the further studies that would need to be conducted if it is to become an effective tool in the study of organizations.

The Application of the Technique

In the analysis that follows, the application of the technique will be broken down into eight steps:

1. negotiating an agreement with the group to be studied;

2. creating the Similarity Questionnaire;

3. collecting the Similarity Ratings;

4. obtaining an organizational chart for the group;

5. creating the data matrix (or matrices);

6. fitting the data to one or more models;

7. analyzing and interpreting the results of fitting the data; and

8. using the findings to interpret and/or intervene;

Steps 1 through 7 are required in any application of the technique. Step 8 will only be used in certain applications (see below).

At each step, there are options in the application of the technique. In many cases, there is little empirical evidence, at this point, to guide the consultant in making these decisions, and some of the further studies that might be conducted in order to provide such guidance will be discussed.

Step 1: Negotiating the Agreement

The successful application of this technique requires the cooperation of as many members of the group as possible, and it seems reasonable to assume that this cooperation can only be secured through an agreement with an individual who has been authorized to speak for the group (or, if necessary, a separate agreement with each member of the

group). The group, for its part, agrees to participate (primarily through filling out the questionnaire). The consultant, on his/her part, must also make commitments on two issues. First, an agreement must be reached concerning the limits of confidentiality for both the responses to the Similarity Questionnaire and for the analyses and interpretations produced by the consultant. The consultant may not think that the confidentiality of the raw similarity data is a sensitive issue, because it seems so unlikely that an outsider could draw any conclusions it. Subjects who are not sure what the purpose of this technique is, or what power it has to reveal their thoughts and attitudes, may wish to be assured that the raw similarity data kept strictly confidential, and this should not present any problems for the consultant.

The confidentiality of the analyzed data (i.e. the configurations that are derived from fitting the raw similarity data to one or more of the algorithms discussed above and the interpretations that the consultant draws from those configurations) is a more troublesome matter. On the one hand, it is quite possible that some of the findings might be painful or even damaging to some members of the group. For example, in both of the studies reported here, individuals were identified who appeared to be having difficulty taking up their role in the group, as indicated

by the fact that their placement in the similarity derived configurations diverged markedly from those that would be expected on the basis of their title and position in the formal organizational structure. Though these problems may have been well known, at some level, to the individuals in question and to other members of the group (as the interview data from the X Team certainly suggests), it may be that the group is not prepared to have this made explicit in public. Furthermore, in a more general sense, there may be other facts about the group that will be brought to light by the configurations which appear to call for change and which the group or its leadership may not prepared to address (for example, the problems with the "Social Work/Psychiatric unit in the X Team).

On the other hand, a consultant who has been asked to propose interventions a group, and who intends to use this sort of data to guide those interventions, may well wish to be able to present the analyses to the group in order to explain and justify the proposed interventions. Therefore, the outcome of the negotiation on the limits of confidentiality will place important limitations on the way in which the data can be used in intervening with the group. In the studies discussed here, none of the findings were reported back to the group, and so the discussion of this issue here should be taken as speculative.

In negotiating an agreement on the distribution of the results, the consultant can distinguish between four levels of disclosure. First, the same full report of the findings can be given to everyone in the group. This will probably constitute a fairly potent form of intervention, putting some potentially difficult the issues out in plain view. If such a course is chosen, one would want to make sure that there was an opportunity for the group to discuss and work through the issues that were raised.

Second, a full report could be given to the group leadership, which is authorized to speak for the group in determining further interventions, but a less than complete report (or no report at all) given to the rest of the group. This choice of this method would seem to suggest that the leadership was mature enough to deal with the issues raised by the data, though the membership was not. It might also imply that the leadership was not directly implicated in the problems that were raised, which seems an unlikely proposition. Furthermore, such a course of action would create a situation in which the leadership knew something that was kept from the rest of the group and in which the consultant had a special alliance with the leadership around these shared secrets. Such a climate might accentuate the suspicion that groups often feel toward consultants and make

the group more difficult to work with.

Third, the report to the group could be edited (as were the hypotheses used in the Second Questionnaire of the Second Study) and a "diluted" version given to the group. Again, this option conveys an implicit message that the consultant knows things about the group that cannot, for some reason, be divulged to the group itself, and one would have to expect that the group members would formulate their own speculations as to why they could not be told. They might, in effect, ask themselves questions such as "what secrets has the consultant discovered?" or "who is the consultant protecting?" A consultant should be prepared for the impact that such an atmosphere may have on his/her ability to work constructively with the group.

Finally, the consultant could give no direct report of the findings to the group, with the understanding that the findings will, however, be used by the consultant in formulating his/her recommendations. This option creates a situation that is essentially similar to that in the third option, but it has the advantage that the group may forget about the similarity questionnaire, or assume that they have not been given the results simply because they are "too technical." As a result, they may be less suspicious.

On balance, it would appear that a consultant would be well advised to choose one of the extremes of the disclosure options discussed above. In other words, one should either put everything on the table and attempt to deal with the consequences (the first option) or simply treat the similarity data as one of many sources of information about the group and focus on the diagnostic formulation and proposals for intervention.

There is one final issue that should be mentioned here. In general, we would expect that the wider the distribution of the analyzed data, the more potent the effect on the group, for better or worse. It would seem, therefore, that any wide distribution would require as comprehensive an informed consent as possible. It will be necessary to let the group know as much as possible about what kinds of findings may be produced by this technique (for example, the fact that it may identify "outcasts"), about the likelihood that these findings are valid, and about the interventions that might be used in the event that difficulties are discovered. Obviously, given the primitive state of current knowledge about the technique, such an informed consent would be difficult. In the studies described here, in which there was no agreement (or intention) to intervene, either no information (the Pilot Study) or very limited information (the Second Study) was provided to the group members. In a

consultation situation, it would probably not be possible to disclose so little information from the findings, at least to the client had retained the consultant.

Step 2: Creating the Similarity Questionnaire

There are several sub-steps in this task.

First, the group of stimulus individuals must be defined. In some cases (the X Team is an example), the members of the group may be defined in a formal sense, and this may present little difficulty. There may well be, however, groups whose membership is a matter of more dispute and it may be more difficult to decide who to include.

There are two questions that bear on this issue. First, how important is it to get the group membership exactly right? Unfortunately, it is conceivable, in the present state of our knowledge, that it is quite important. If we accept Tversky's contrast model, we assume that there is an initial stage in which the person making the similarity judgment selects a few features (out of the very large number of features possessed by the stimulus objects) to be used as the criteria for the similarity judgments. In the procedure followed in the studies described here, this decision is made after the subject has been given a broad

hint as to what the set of individuals will be. (The similarity question was posed after the subjects have been asked to indicate how well they know the people in the group.) It possible, therefore, that the choice of the features is influenced by the specific set of individuals that are to be compared. If this were the case, then the use of sets of stimulus individuals that differed in composition might result in quite different similarity ratings and, therefore, quite different output configurations. Again, when the membership of the group is unambiguous, this will not be a problem, but this may not always be the case.

There is another situation (even more common, perhaps) that may be encountered in the application of this technique that raises a closely related difficulty. There may also be instances in which the group to be studied contains so many individuals that the number of all possible paired comparisons (N*[N-1]/2) becomes unmanageable. For instance, the 25 stimulus individuals used in the Pilot Study resulted in 300 comparisons and the Similarity Questionnaire took most of the subjects approximately 35-45 minutes to complete. This is probably close to the upper limit of what subjects will tolerate, but many naturally-occurring groups will have more than 25 members.

In this situation, there would appear to be two options: 1) one can eliminate some of the stimulus individuals, reducing N and therefore reducing [N*(N-1)/2]; or 2) use the full set of stimulus individuals but discard some of the paired comparisons and, therefore, reducing [N*(N-1)/2] directly.

In the first method (Method 1), one would select as representative a subset as possible of the group membership and use this group as the stimulus individuals. In choosing such a subset, it would appear to be prudent to include all of the leaders of distinct sub-units of the group (i.e. the top layers of the organizational chart) and fill out the rest of the set of stimulus individuals with approximately equal numbers of the lower-level staff from each of the sub-units. (This is, in fact, what was done in the Pilot Study.)

In the second method, one would use the full membership of the group as stimulus individuals, but discard some of the paired comparisons, reducing the total number of comparisons in the questionnaire down to a manageable number. The justification for this would be that there is a very considerable degree of redundancy in the use of the full set of [N*(N-1)/2] paired comparisons, and that some of these estimates could be discarded with no serious loss of

information. The risk would be that by reducing the number of estimates of psychological distance, one loses the redundancy that corrects for error and inconsistency in the estimates, and that the proportion of error variance in the resulting data matrix will grow too large.

The follow-up study described below could test and compare the effectiveness of these two methods.

Follow-up Study #1:

This study could be carried out using a group that was clearly too large to ask subjects to rate all possible pairs of stimulus individuals (for example, 42). A general form of the organizational chart is given as Figure 16. There is an overall leader for the group, which is then divided into three functional sub-units (A, B, and C), each with a leader (L_A, L_B, and L_C), and number of staff (the numbers given are solely for the purpose of giving examples). Four versions of the Similarity Questionnaire would then be prepared:

Two versions would be prepared according to first method described above. Both versions of the questionnaire would include L, L_A, L_B, and L_C, but the staff of each of A, B, and C would be randomly split in

Figure 16:
A Generalized Organizational Chart

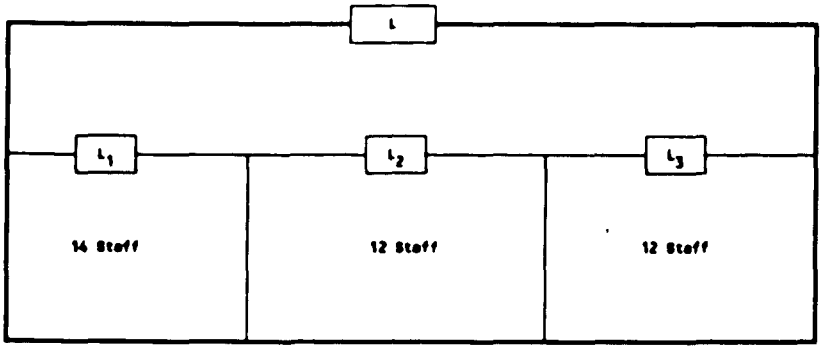

half and different halves would be used to fill out the set of stimulus individuals. The result would be two different Similarity Questionnaires, each with 253 comparisons.

Two questionnaires would be prepared according to the second method described above. The full set of 42 stimulus individuals would be used, but the 861 possible paired comparisons would be randomly split into three different groups, each with 287 comparisons. Two of these sets would be used as the items for the two different Similarity Questionnaires.

The subjects would be randomly given one of these questionnaires, and the resulting data would be analyzed according to the rules of inference discussed above. If the configurations derived from the two versions of a single method differed, it would appear that the method was invalid. If the configurations derived from the two different versions of the same method were similar to each other, but differed from those derived from the other method, one would need to test the validity of the inferences that could be drawn from the resulting output configurations through a method that was equivalent to the Second Questionnaire procedure described above, to see which, if any, of these different methods appeared to be

capable of producing valid inferences.

Groups with small numbers of stimulus individuals present a different problem. With a group of 10, for example, the number of possible pairs (45) is relatively small in relation to the number of stimulus individuals. One might be concerned, in such a case, that the number of possible configurations is relatively large while the number of constraints on the configurations, represented by the number of similarity judgments, is relatively small. If the technique were to be applied to small groups, some additional studies designed to check the validity of the inferences drawn, should be performed.

For the remainder of this discussion, I will assume that the number of stimulus individuals is approximately the same as was used in the studies described above and that none of these special problems have arisen.

Once the group of stimulus individuals has been defined, a Similarity Questionnaire must be constructed. An issue that arises here is the question of the order in which the stimulus pairs are presented. As noted above, the studies described here employed a "Ross-ordering" (Ross, 1934) in order to be consistent with the methods employed in the Jones and Young study. It is possible that other orders

could be employed. Different orders could be compared in the manner suggested above for comparing different versions of the Similarity Questionnaire. One easy study would be to compare the results obtained with a Ross-ordering and a reversed Ross-ordering. It would be hoped that no difference would be found and, if there was, it could throw significant doubt on the utility of the method.

The final step in preparing the Similarity Questionnaire is to decide on the wording of the "similarity question." In the two studies described here, the very open-ended form of the question used by Jones and Young was employed. In part, this was done to ensure that the results obtained could be considered to be directly comparable to Jones and Young's results. Beyond this, however, there is a more substantive issue. What are the arguments for and against the open-ended form of the similarity question?

The open-ended form of the question forces the subject to make his/her own decision as to the criteria for similarity (in other words, to formulate an answer to the question "similar with respect to what?"). In the terms of Tversky's contrast model, it allows the subject the freest choice as to what features will be taken into account in making the similarity judgment. Most important, it appears to constitute the best, and most stringent, test of the

theory of sentience advanced by Miller and Rice.

Their central contention was that effectiveness in task performance is maximized when the members are strongly identified with the sub-group to which they are assigned- that is, when task and sentience boundaries coincide. In this circumstance, the members are most strongly committed to their task. It is also possible, however, that task-irrelevant identifications may intrude into the work situation and interfere with commitment to task. This technique was intended as a way to determine the distribution of sentience across the different possible groups within the organization, and to assess the degree to which task-irrelevant identifications intruded into the group's mental life. Given this purpose, it seemed wise to frame the similarity question in the most open-ended way possible. Though it would be easy to write a similarity question that would virtually guarantee that task relevant groups would emerge from the data (for example, "indicate the degree to which the following individuals are similar with respect to their work roles"), or that a form could be found that encouraged the emergence of non-task relevant categorizations (for example "indicate the degree to which the following individuals are similar with respect to their personalities"), it is not clear what conclusions about the organization under study could be drawn from the responses

to such questions. The responses to the open-ended form appear to have the virtue that it makes it possible to estimate the balance between the different types of identification.

The assumption here is that subjects are, as they judge similarity, striking a balance between two different sets of similarity relationships, and that the degree to which the balance is tipped in one direction or the other has an implication for the effectiveness of the organization. This interesting speculation could be pursued in a follow-up study of the type described below.

Follow-up Study #2:

Suppose, for a moment, that there are, in fact, two major types of categories that can have an impact on similarity: 1) task-related categories; and 2) personality categories. We could then design two different forms of the similarity question, each intended to elicit the set of similarity ratings that are determined purely by one or the other type groupings, construct two different Similarity Questionnaires based on these forms of the question, and, presumably, obtain a different output configuration for each set of categories. One could also obtain a set of responses to the open-ended form of the question.

For each pair of stimulus individuals, we would have three ratings: 1) a rating in response to the open-ended question (Unstructured Similarity, or S_U); 2) a rating in response to the task-directed question (Task Similarity, or S_T); and 3) a rating in response to the personality-directed question (Personality Similarity, or S_P). The question could then be asked, can the open-ended similarity ratings be regarded as a weighted composite of the task and personality ratings? In other words, does:

$$S_U = W_1S_T+W_2S_P$$

where W_1 and W_2 are weights which can determined by linear regression?

If it should prove to be the case that Unstructured Similarity can be treated as a function of Task and Personality Similarity, then one might be inclined to regard configurations produced on the basis of the Unstructured Similarity ratings, such as those used in the studies described above, not as "the map of the social space" but as the result of the superimposition of a "task map" and "personality map." It also raises the possibility that individuals might differ in the weights that they assigned to each type of similarity in coming to an estimate of

Unstructured Similarity[21]. In other words, there may not be one map, but two maps, both of which are drawn upon and the results combined when Unstructured Similarity is rated.

Step 3: Distributing the Similarity Questionnaire and Collecting the Data

This task is relatively straightforward, but there are two issues that are worthy of brief mention. The first is the confidentiality of the questionnaires. The consultant will, most likely, wish to know which individuals gave which responses, but it is possible that at least some of the subjects will resist this. At this point in the development of the technique it is difficult to say how important it really is to know who the individual respondents are. This is particularly true because the data-fitting procedures that have proved to be the most useful, ADDTREE and EXTREE, do not have any capacity to take individual differences into account. As the technique and the model-fitting procedures evolve, strong reasons for preferring to know who the subjects are may well emerge.

---

[21] Such individual differences would be similar to the model of individual differences employed by Carroll and Chang in their development of INDSCAL (Carroll and Chang, 1970).

The second point concerns the individuals who do not complete the questionnaire. In both of the studies reported here, a few individuals never completed the Similarity Questionnaire. My impression is that the individuals who were later tentatively identified as "outcasts" were somewhat more likely to be among the non-responders than "non-outcasts." If this impression were borne out in subsequent uses of the technique, it might suggest that refusal to cooperate might be regarded as a very "soft" sign of "outcast" status. If this were the case, there are two speculations that one might advance to explain the finding. First, since the use of the technique is negotiated between an investigator (or consultant) and the group (even if the investigator deals with each individual member), failure to comply represents an unwillingness to commit to a course of action that has been taken by the group. Such an unwillingness seems to be very consistent with "outcast" status. Second, it is possible that "outcasts", if they are, in fact, individuals who are having difficulty in taking up their role in the group, are in more emotional distress, at least with respect to their work situation, and that they will be concerned that feelings that they struggle to control in the course of their working life will be revealed through the technique.

Step 4: Obtaining an Organizational Chart

In the studies reported here, the organizational chart (the formal description, or theory, of the task structure) were taken from single sources who were presumed to be authoritative on the subject. In many organizations, the organizational chart may be well-defined and a matter of official record (as it was in the academic health center) but in others, there may be a much less well-defined task structure (as was, perhaps, the case in the interdisciplinary team). In such situations, it might be advisable to collect a number of charts from different individuals.

Step 5: Creating the Data Matrix(ices)

Once the Similarity Questionnaires are received, the investigator has full sets of similarity ratings from each of the respondents. Each of these sets can be used to create a triangular proximity matrix of the type that is used as the input for all of the model-fitting procedures used in the studies described above. If these individual data sets are analyzed in this way, configurations will be produced which could be thought of as each individual's map of the group. Another approach, which was used in the studies described above, is to average the ratings for each

pair of stimulus objects, creating a matrix of averaged estimates, and to then use this estimate as the input for the data-fitting procedures.

In principle, it would appear that both approaches could be valuable and the method of individual maps might well be preferred. After all, the point of the procedure is to allow the investigator to reconstruct the "map in the head", and each of the individual data sets comes from one person's head, while the averaged data set comes from no one person's head. It might well be the case, in a group of individuals whose ratings were quite discrepant, that the averaged ratings corresponded to no individual's ratings. If that were the case, it would seem unlikely that the post-analysis configurations could tell the investigator anything useful about the functioning of the group. The individual ratings, on the other hand, would seem to be always relevant, because they allow the investigator to reconstruct each individual's "map."

Surprisingly enough, though, the indications from the research described here are that the averaged data is quite useful (at least in these two groups, at these times) but that individual data may be much less useful than might, at first, be thought. The individual data sets from the respondents in the Second Study were fitted to the ADDTREE

and EXTREE models, producing individual EXTREE and ADDTREE configurations for each individual, but a number of difficulties emerged.

The first is a practical one, in that the investigator in confronted with a large number of configurations and it is difficult to know how one can set about interpreting them. In this case, it appeared that though there were some structural similarities, there were many structural features that seemed to be unique to a single individual's configuration. Furthermore, it appeared to the investigator that the configurations diverged very markedly from the configuration that would have been expected from the organization chart. Indeed, the configurations seemed to be basically different from the organization chart, and different in ways that were not immediately interpretable, at least in the light of the very limited information available to the investigator. This is a considerable contrast to the configuration produced from the averaged data, which was much more easily reconciled with the organizational chart, and where the deviations from the chart lent themselves quite readily to interpretation.

It is possible, however, that the congeniality of the averaged data is misleading and that the individual data should still be preferred (on the grounds discussed above),

even if it is less easy on the investigator. Against this notion, however, one might offer the following hypothesis.

One of the assumptions that underlies the technique employed here is that subjects employ the same version of the "similarity function" (or the same "map in the head") in all of the N*(N-1)/2 comparisons that they make as they complete the lengthy Similarity Questionnaires. In the terms of the contrast model, we assume that subjects stick to their initial choice of relevant features, that they stick to their initial decision about the relative importance of common and distinctive features, and that they stick to their initial decisions about the relative weight of each selected features. In a sense, it is as if we assumed that the subjects were using an absolutely rigid yardstick to measure a long series of distances. On the other hand, suppose that subjects are using yardsticks that are quite flexible and inclined to warp and stretch as they are carried from one place to another. In this case, we would expect that there would be a considerable amount of random variation introduced into the measurements made by the subjects.

If something of that sort were to occur in the sort of procedure that we have been considering here, we might expect that the estimates of similarity obtained from each

individual's questionnaire would contain a considerable amount of inconsistency, which is, for our purposes, the equivalent of meaningless "error" variance. When we then fit the data from an individual questionnaire to the various models, we can expect that the configuration that results will have been shaped so as to take account of that meaningless variance and will, to that extent, be distorted and uninterpretable. In fact, if the proportion of meaningless variance to total variance is high, we could expect that the "true" configuration might be submerged by the "meaningless" variance.

In such a circumstance, averaged data may appear more attractive. With averaged data, we might hope that the influence of such "error" variance will have been greatly diminished by the averaging, which will have emphasized those aspects of the judgments which are consistent, both within and across subjects. As a result, we might hope, the output configurations from averaged data might be more likely to be interpretable.

Some informal and anecdotal information can be offered to support this hypothesis. Though the issue was not explored systematically in the studies described above, several subjects volunteered that they were aware of having varied the criteria that they used in making the similarity

judgments.

This issue could be explored further in a study of the following type:

Follow-up Study #3:

The design of this study would be generally similar to the design of the Second Study, with two exceptions. First, both averaged data and individual data would be fitted to the EXTREE model, which would result in an output configuration for each individual and an output configuration for the averaged data. All of these configurations would be analyzed, according the procedures described here, and hypotheses concerning the group would be generated. In addition, a set of test hypotheses, not supported by any of the configurations or the organizational chart, would be created. The result would be a set of four different types of hypotheses that could be included in, and tested by, a Second Questionnaire which can be given to each individual subject:

1. Hypotheses which are supported by both the averaged data configuration and the individual's configuration;

2. Hypotheses which are supported by neither the individual nor the averaged data configurations;

3. Hypotheses which are supported by the averaged data configuration but not by the individual data configuration; and

4. Hypotheses which are supported by the individual data but not by the averaged data.

One would expect that the first class of hypotheses would have the highest average agreement scores and that the second would have the lowest. The more interesting comparison would be between the third and fourth classes. If the individual agreement scores for the third class were higher than those for the fourth, the use of averaged data would be encouraged and the line of reasoning outlined above would be supported. If the reverse were true, it would appear that individual data would be preferred in future applications.

In the event that averaged data is preferred and used, some precautions might be in order, to make sure that the data is coherent enough to be appropriate for averaging. For example, it is conceivable that a single group might include two distinct subgroups, each of which was inclined

to approach the similarity judgment task in distinctly different ways. If this were the case, there might be, in effect, two different sets of judgments, which bore very little similarity to each other. The result of averaging the individual data might be to create a data matrix which bore little resemblance to the matrices for either subgroup, and an output configuration that was quite irrelevant to the psychological life of the group.

One way to test this would be to carry out a "Points of View Analysis" (Tucker and Messick, 1963). In this procedure, each subject's ratings (for all pairs of stimulus objects) is correlated with the ratings of every other rater. The result is a set of correlation coefficients which express the degree of similarity between the ratings of all possible pairs of raters. These correlations can be treated as a sort of proximity measure, and the set can be organized into a data matrix that can be fitted to multidimensional scaling (or one of the other models discussed here). Inspection of the resulting output may give an indication as to the degree of consistency of ratings within the group. If, for example, the output suggested that there were two very distinct subgroups, the members of which had ratings that were very similar to each other but very different from those of the other subgroup, one would probably be wise to separate the two groups and

average and analyze their data separately. Such a finding would, of course, be extremely interesting and, most likely, significant with respect to the psychological life of the group.

Step 6: Fitting the Data to Models

Once the input data has been collected, the investigator must decide which of the many different algorithms that are available for analyzing proximity data he/she wishes to use. All of the algorithms described above accept the same sort of input, and so all can be used simultaneously. At this point, however, one can begin to formulate some general conclusions about the advantages and disadvantages of the different models. The discussion will focus on the multi-dimensional scaling, additive tree, and extended tree models.

Multidimensional scaling is the longest established of the various techniques and has the most extensive body of applications behind it. It would appear, however, on the basis of the studies reported here, that it has relatively limited application to the psychological life of groups. This may not be very surprising, if we recall Tversky's assessment that spatial models such as multidimensional scaling are best suited to stimuli and situations in which

the relevant features are relatively few and where those features vary continuously. The classic examples are perceptual stimuli such as tones or colors. For stimuli in which there may be a larger number of relevant features, and where those features have an "all-or-nothing" quality (e.g. gender) the spatial model may be far less appropriate, for several reasons.

The first is practical. Spatial models can only take account of independent features through the addition of an orthogonal spatial dimension to the model. In multidimensional scaling, the number of dimensions to be employed is specified by the user and the optimal number is not easily determined. Even more important, configurations in three dimensions are already somewhat difficult to interpret or visualize, and configurations in more than three dimensions are extremely difficult to interpret.

The second is more subtle. As noted above, multidimensional scaling accounts for the similarity variance by adjusting the placement of the stimulus objects on a bipolar continuum which can be thought of as representing a feature. It has, therefore, the capacity to account for a great deal of variance through the use of a relatively small number of features. In instances where one suspects that the relevant features are likely to be more

numerous (and more "all-or-nothing") one is not likely to obtain a satisfying configuration from a model whose fundamental assumptions are different.

Despite this, however, multidimensional scaling has (at least) one signal virtue, which is its ability to represent the "centrality" of group and subgroup leaders. In both of the studies reported here, the multidimensional scaling output placed the overall leader in what could, after examination, be seen to be a central position, equidistant between clusters representing the major subgroups of the organization. This would appear, intuitively at least, to be an appropriate placement, and it is one that cannot be accomplished by the tree models, which place all individual stimuli at the end of the branches of the tree. In the tree models, one must infer centrality on the basis of characteristics such a length of extension which are much less compelling.

In groups with a relatively simple structure, such as the academic health center discussed above, this sort of analysis may produce some useful information about the organization. Success requires a certain degree of luck, however, to the degree that one needs to have happened on a group whose structure lends itself to the multidimensional scaling representation. In the academic health center

group, for example, it seems that there were only four major functional clusters, a number which is small enough that their interrelationships can be well expressed by placing them at the corners of a rough square in a two dimensional space. In the interdisciplinary health care group, it appeared that the true number of subgroups was somewhat greater, and that the subgroup structure could not be easily represented in a two dimensional space. As a result, the output configuration appears to have collapsed a large number of subgroups together into a small space, obscuring, rather than revealing, their interrelationships. Further applications will shed more light on this question, but at this point, it seems likely that this problem will prove to be a serious practical limitation on the use of multidimensional scaling.

The additive tree and extended tree models appear to be much better suited to this application. This is, no doubt, due at least in part, to the fact that its assumptions (large numbers of relevant features, "all-or-nothing" membership in a category) are more appropriate to the nature of the stimulus objects and the representations involved. In addition, however, the tree models have a significant practical advantage. They are easily interpretable, especially in comparison with higher dimensional spatial configurations. The tree models can display a complex

structure in an easily accessible way, and they appear to be less dependent on the group under study having just the right sort of structure- that is, they appear to be more robust and to be capable of application a wider range of groups than the scaling model.

A potential disadvantage of the tree models is that they do not, at this time, have the capacity to take individual differences into account in the way that INDSCAL can. At this point, however, it is not clear how individual differences are to be made use of in this sort of application. Furthermore, it seems likely that individual differences versions of the tree models will be developed, or could be developed if it appeared that there was a need for them. One compensation offered by the EXTREE algorithm is the fact that it allows the user to test the salience of a particular grouping (e.g. gender, race) by measuring the increase in variance accounted for when the cluster is added to the model.

In summary, two points are worth reiterating. First, all of the algorithms discussed in these studies accept the same type of data as input. The investigator is, therefore, free to use any and all of the techniques that he/she wishes- to, in a manner of speaking, reprocess the ore with one technique after another until all possible meaning has

been extracted from it. Second, it appears that certain types of group structure (or certain features of group structure) lend themselves to being revealed by certain techniques. If this is the case, it would suggest that multiple techniques should be used together in the attempt to extract information from the similarity data.

Step 7: Interpreting the Output

Once the data has been fitted to the various models, the configurations that result must be interpreted. A first step should be the examination of the variance accounted for statistics for each of the output configurations, which may give some clues as to the overall goodness-of-fit of the configurations. At this point in the development of this technique, there are no standards for goodness-of-fit, but it would appear, on the basis of these studies, that one should expect to be able to account for 65%-90% of the total similarity variance.

Much of the discussion here is a recapitulation of the processes of analysis that have been described in the report of the results of the two studies. As many of the rules of inference apply to a specific algorithm and its output, the two major techniques, multidimensional scaling and the tree models, will be discussed separately.

Multidimensional Scaling

The use of this technique involves a preliminary step which is not demanded by the tree models- that is, deciding on the dimensionality of the space that in which the algorithm will construct the configuration. Multidimensional scaling algorithms permit the investigator to specify the number of dimensions that will be used. In making this choice, the investigator has to balance the increase in variance accounted for that can be expected as the number of dimensions increases against the very severe difficulties that are encountered when one attempts to interpret arrangements of points in spaces of more than three dimensions. The most straightforward approach, which is probably the best at this point in the development of the technique, is to obtain configurations in two, three, four, and five dimensions and calculate and graph the increments in variance accounted for as the number of dimensions increases. If the increments obtained by using the higher dimensions are small, and the graph of the increments is flattens out, one is probably better off with the two or three dimensional configurations. If, on the other hand, there are some substantial increments in variance with higher dimensions and, in particular, if the graph of the increments moves up after having levelled off, it suggests

that the higher number of dimensions has allowed the algorithm to capture an important aspect of the group structure, and one may have to struggle with the difficulties of interpreting the high-dimensional spatial arrangement. The discussion below will assume that the three dimensional configuration has been chosen, but the same principles can be applied to configurations of any dimensionality.

Once the number of dimensions has been chosen, the next step is to locate the overall leader of the group, if one can be identified, and to see if that person is: 1) near the extreme of one dimension, in which case that dimension may tentatively be identified as "status in the hierarchy"; and 2) at the center of the plane defined by the (other) two dimensions, in which case these dimensions may be tentatively identified as the space in which the functional clusters will be located. After this, one should attempt to locate the subgroup leaders, who may occupy similar "above and in the middle" locations in relation to the members of their subgroups.

After the locations of some of the key figures has been determined, one can begin to compare the MDS configuration with the structure depicted by the organizational chart. The comparison should focus on both the location of

individuals and the identification of subgroups (though the latter task is not as easily done in multidimensional scaling outputs as with the tree models).

The starting point for the comparison could be the assumption that what the MDS configuration represents is a translation of the structure depicted in the organizational chart into the three dimensional space. Therefore, using the overall leader and subgroup leaders as landmarks, the investigator should attempt to locate the other members of the subgroups with which they are associated. (As indicated above, if there are more than four substantial subgroups, and one is using a three dimensional configuration, it is likely that some of these subgroups, most likely the smallest ones, will have been superimposed on top of each other and the investigator will have to disentangle them in his/her imagination.) Several types of discrepancies should be searched for. The first is subgroup leaders who appear to be lower on the status dimension than their assigned roles would suggest. A second is individuals who are placed at some remove from the other members of their task-structure-assigned subgroup. If a given subgroup contains only one such individual, it is most likely an indication that that particular individual is having difficulty taking up their role. If there are a number of such individuals, it suggests that the subgroup in question lacks coherence or

"psychological reality."

The Additive and Extended Tree Models

The tree models, ADDTREE and EXTREE, are more convenient in that they do not require the investigator to do anything equivalent to specifying the dimension of the configuration. The algorithms produce a readily intelligible configuration directly from the similarity data. The tree models also differ from multidimensional scaling in that they make the structure of subgroups explicit, and the investigator does not need to work to extract it.

One of the difficulties presented by the configurations from the tree models, in fact, is the wealth of detail in the tree structure. The investigator is confronted with the N stimulus objects, each with its own unique extension, with what may be a large number (as many as N-2) of nested subclusters, each with its own extension, and with (in the EXTREE configuration) a set of superimposed cross-cutting clusters, each with its own weight. In order to make sense of this information, both quantitative and qualitative tests need to be applied.

First, the quantitative tests. The extension lengths

provided by the tree outputs for all of the subclusters and for each of the individual stimulus objects can be taken as indications of the distinctiveness of each of the clusters and individuals. One of the first questions that the investigator confronts is the task of deciding which of these statistics have some significance. In the studies reported above, it appeared that one way to do this is to calculate the average extension lengths for each group of extensions (i.e. subclusters and individuals) and the z-score for each subcluster and individual (with respect to the appropriate group mean). If the findings reported above were to prove generally applicable, one could make the following tentative conclusions: 1) individuals with exceptionally large extensions (z-score > 1) may prove to be individuals who are experiencing difficulty in taking up their roles; 2) individuals with exceptionally small weights (z-score < 1) may be leaders of the subgroups, with a stronger than average bond to the overall group; and 3) subclusters with exceptionally large weights are likely to correspond to salient ("psychologically real") subgroups in the organization.

The qualitative task is very similar to the that described in connection with the multidimensional scaling model, that is, the comparison of the structure depicted in the tree configuration with the ideal structure depicted by

the organizational chart. The specific questions are essentially the same, but the tree models are particularly well suited to explore the degree to which the subgroups that are identified by the organizational chart are, or are not, "psychologically real" for the members of the organization. One useful procedure is to attempt to find a match for each organizational chart subgroup with a cluster in the tree output.

The investigator should also search for individuals with anomalous locations. In some cases, individuals will be found who are bound together to other individuals with whom they would appear to have no close association in the task structure. In such instances, further investigation may be needed to determine whether these bonds reflect personal, non-task related factors, or genuine task relationships not depicted in the organizational chart.

There may also be instances where individuals appear, in the tree configuration, to be isolates, with no association to any other members of the group, or where such connections are exceptionally weak. Such individuals may prove to be persons who are having difficulty taking up their role in the group.

Identifying the leadership in a group is somewhat more

difficult in the tree configurations than it is in the scaling output, but it would appear, on the basis of the studies reported above, that there are several places in which one can look. First, as noted above, there is a good possibility that leaders will be found among the individuals with exceptionally short extensions. The rationale for this supposition is that subgroup leaders, though they have strong bonds to the other members of their subgroup, also have unusually strong bonds to the other subgroups of the organization and to the leader, and that these bonds are expressed by a lessening of the individual's overall distinctiveness. In other words, a Vice President has a strong identification with the sector for which he/she is responsible, but also has a strong identification with the other Vice Presidents and with the President. Another possible clue to subgroup leader status is the location of the individual at the center of the subgroup, reflecting the degree to which the subgroup leader should, ideally, be equally identified with all members of the subgroup. A final clue to the identity of the leadership group can be the emergence, in the EXTREE output, of a cross-cutting cluster that contains the overall leader and the subgroup leaders. Such a subgroup (the "inner circle") may be where the real executive authority resides and may differ in composition from the group that one would expect based on the organizational chart. (The academic health center

studied in the Pilot Study provides good examples of all of these indicators.)

As mentioned above, the EXTREE algorithm also allows the investigator to test the salience, or psychological significance, of specific non-task related groupings such as gender or race. The is done, essentially, by instructing the computer to add a given cluster of individuals to the model arrived at by the ADDTREE algorithm, and seeing if there is a large increase in variance accourted for. If there is, it supports the hypothesis that this is a psychologically significant subgroup. If the increment is small, however, the converse cannot necessarily be concluded. It may well be that the variance associated with the subgroup has already been accounted for in the weights assigned to some of the clusters in the ADDTREE configuration, leaving little residual variance to be accounted for by the added cross-cutting cluster. One way to assess this possibility is to reexamine the ADDTREE configuration to see if the members of this subgroup appear to be concentrated in particular subclusters or are spread evenly across the configuration. Only if they are spread evenly can one begin to feel some confidence that the group is not a psychologically significant one. Caution in interpreting this "test" seems particularly appropriate in view of the fact that neither of the two studies reported

here found any significant increases in variance accounted for with any of the most obvious non-task related subgroups, suggesting, to this investigator at least, that the proposed "test" is too insensitive to be useful. Further studies, particularly in groups where these groupings are known to be psychologically significant, would be very useful in assessing this particular aspect of the technique.

Global Interpretation

After analyses of the type described above have been conducted on the output configurations from whatever algorithms have been selected, the investigator (or better, the consultant) may wish to discern the meaning of the whole pattern that is depicted by the output configurations. This section, which will discuss such interpretations and use the two studies reported here as examples, is necessarily more speculative than the sections above, in which most of the findings have been confirmed by the data from the Second Questionnaire and the Interviews. For the most part, the hypotheses discussed below have, at best, only the very weakest, most tentative support from the data. The are intended as examples the way in which this sort of data, taken together with all of the other information that is available on the group and with the investigator's "clinical judgment", can be used to generate hypotheses that are more

comprehensive.

In formulating a "global interpretation," the investigator or consultant is attempting to make sense of the pattern of sentience presented by the output configurations (including all of the sorts of findings discussed above- the levels of organization that are, or are not, psychologically real, the "outcasts", the locations and identities of the leaders, etc.) in a the context of what is known about the organization's task and the technological and social psychological constraints under which the task is being executed.

Of the two groups studied here, the interdisciplinary health care team probably provides the richest and most intriguing example. In this group, individuals bring to their roles very different experiences and training, and those roles confer very different responsibilities and privileges. Some members of the Team, for example, are members of a medical school faculty, with the implication that they are expected to maintain a degree of academic productivity (e.g. publications, research grants, etc.) and that they have a relatively high degree of professional autonomy. Other members of the Team are given more conventional professional health care provider roles. There is a considerable difference in the pay provided to the

different members of the Team, and different assumptions about what sorts of accomplishment will lead to what sorts of professional advancement.

The primary task of the Team is to provide medical and social services to individuals and families who have been infected with a frightening, thus far invariably fatal, disease.

As one looks at the EXTREE output (Figure 12, page 125), the most striking global feature is the fact that there is only one large cluster (the "main cluster"), containing seven of the eighteen stimulus individuals, while the rest of the individuals (the "satellites") appear as isolated individuals or in rather tightly bound pairs, loosely linked to other pairs or to the larger group. The main cluster, furthermore, contains all of the physicians except the two physicians who have been tentatively identified as "outcasts." Four of the seven members of the main cluster are physicians, and it appears, on the basis of the interview data, that the physicians on the Team have a more limited involvement with the patients treated by the Team. The remaining members include the Secretary (a non-patient care position), a Physician's Assistant who is described in the interviews as having a de facto administrative position and is one of the founding members

of the Team, and a Nurse Practitioner who is moving out of clinical work into the management of a research project associated with the Team.

Apart from the two "outcast" physicians, the "satellites" consist of the four social workers, the two nurse educators, two Physician's Assistants, and the Phlebotomist (also tentatively identified as an "outcast"). The interview data suggests that the psychological burden of the care for these severely ill patients may fall most heavily on these individuals.

The image that came to the mind of this investigator was of the situation (or, at least, one view of the situation) on the front lines in World War I. Behind the lines, there are the headquarters of the commanders, safe from immediate danger, and comfortably lodged in a French chateau. On the lines, exposed to direct fire and living in an intensely hostile and life threatening environment, are the front line troops. For protection, they dig themselves into the ground and rely totally on their companions in those trenches.

In the context of the X Team, I am inclined to see the "main cluster" as the generals and the tight pairings that are characteristic of the "satellites" as corresponding to

very close personal bonds between persons with similar, if not identical, work roles. A primary function of those personal bonds might be to provide a sense of connection and companionship in the face of a bleak and dangerous environment. In comparison with these intense pairings, higher levels of organizational structure (e.g. the Social Work/Psychiatric Unit envisioned in the organizational chart) have no "psychological reality", in the way that, perhaps, larger levels of military organization had a less palpable psychological presence than did the soldier's immediate companions. They may be recognized to exist, but they provide little immediate comfort or protection against the immediate pressures.

Obviously, this is just an example of the sort of inference that one could draw from this type of data. The point is not that it is necessarily correct, but to suggest that such inferences can be the starting point for fruitful discussions with the members of the group concerning the difficulties of their task and roles and of the organizational structure that has been constructed to accomplish that task.

Step 8: Using the Findings to Interpret and/or Intervene

The final task that presents itself to the consultant,

if not to the investigator, is decision as to the use, if any, that will be made of the findings that emerge from the sort of analysis described above. This is a topic on which the studies reported here can have little to say, as no attempts to report or intervene were made, but, in the final analysis, the practical utility of the technique described here will depend on the degree to which it can be used to enable groups to function better. This section, therefore, will attempt to identify some of the options that a consultant has, and will suggest some ways in which they might be explored and compared in follow up studies.

Some of the basic issues have been discussed briefly above in the section on confidentiality and the distribution of the results. In any practical use of this technique, the consultant will be obliged to make a commitment to the group members concerning the use to which the results will be put. The first choices that the consultant and the group will make will, therefore, be between the different possible ways in which the findings will be distributed. There appear to be four different approaches to handling the distribution of the information.

First, the consultant can agree to present the full set of data and findings to the group, either through distributing a written report, or by giving an oral

briefing, or both. This is the approach that a researcher might well prefer to take, as it provides a clear and direct test of the validity of the findings, but it is not clear that it would be helpful to the group. Furthermore, as noted above, it might prove very damaging to some members of the group. In general, a useful rule of thumb might be that this technique is not likely to bring to light facts that have been entirely unknown to the group members, but it may well identify issues that the group has decided, in effect, that it does not feel able to discuss directly. If the consultant brings them forward, the group could react in many ways- by rejecting the findings, by attempting to place the blame for the problems, etc.

A second possibility, that may be more tempting, is to fully reveal the findings to the group leadership, presuming that the group was willing to authorize this. The assumption that is implicit in choosing this course, over the "full disclosure to the full group" option, is that though the group may not be mature enough to cope with the findings, the leadership is. There may well be cases in which this is true, but in those cases where one cannot be so sure, the objections raised above with reference to the whole group may apply equally to the group leadership. They may be just as resistant as the other group members, or even more so if they feel that it was their responsibility to

prevent any problems from having arisen.

The third possibility is that some partial report of the findings be made to the group or its leadership. It seems likely that this will often be the case, and that the precise nature of the findings reported will depend on the plans for future work with the group.

Finally, there is the possibility that the consultant will use the findings from the analysis of the similarity data in conjunction with other sorts of data (interviews, information on the performance of the various units, etc.) in formulating a more general statement of the situation of the organization and of its difficulties and a set of interventions that might be recommended in order to address these difficulties.

The decisions that are made on these issues will depend, to a considerable degree, on the plans, if any, for intervention with the group. Intervention is an issue, as noted above, on which the studies reported here have nothing new to contribute, but the findings, in some instances, suggested specific areas or problems in which intervention might be useful. In the case of the interdisciplinary health care team, for example, it would seem that a consultant would have to consider some of the following

possible areas for intervention.

Intervention #1: The Problem of the "Outcasts"

As noted above, one of the most striking findings in these studies has been the ability of this technique to identify individuals in the group who are having difficulty in taking up their roles. However, the fact that an individual has been so identified does not necessarily indicate what the nature of the difficulty is and there are many possibilities. In some cases, the person may be poorly suited to the role, and lack the experience or capacity to carry it out. In other instances, the group may have unreasonable expectations of the role and it may be impossible for anyone to make a success of it. In yet other instances, it may be that the individual's role places them in conflict with the rest of the group, and that the sort of isolated position that has been considered, in the discussion above, to be an indication of "outcast" status simply reflects this fact. In short, the identification of a given individual as a possible "outcast" demands that the consultant investigate the issue further to see what organizational fact underlies the finding and to determine if an intervention may be required. In some cases, it may be necessary to help the person in question find another role, in others it may be more appropriate to redefine the

nature of the role itself.

Intervention #2: Making the Social Work/Psychiatric Unit Psychologically Real

In the Second Study, the similarity data, the Second Questionnaire responses, and the Interviews consistently indicated that the Social Work/Psychiatric sub-unit envisioned by the Medical Director in his description of the task structure of the X Team was not a "psychologically real" working sub-unit in the minds of the Team members. A finding such as this- that there appears to be a substantial discrepancy between the functional divisions prescribed by task structure and what we may tentatively infer to be the actual distribution of sentience within the group- poses an obvious target for intervention, especially for a consultant working in the Tavistock tradition.

In approaching this problem, one could begin with the assumption that the members of the putative sub-unit do not think of themselves as a group committed to a common task and that it is quite possible, therefore, that this task is not being performed. A first step, then, will be to consult with the Team leadership and the members of the Team to determine what the task is, why it is important, and whether it is, in fact, being performed adequately. It may turn out

that the proposed task is not particularly important, and that there is no need for the putative group. In this case, the result of the intervention will be a revision of the task structure (i.e. the organizational chart), which is, after all, only a theory about how the group has to be organized in order to get the work done. Theories, even the leadership's theories, can be wrong.

A second possibility is that there is an important task for this particular sub-unit of the Team and that it is not being performed adequately and, in such a case, an intervention may required. In a general way, the purpose of the intervention will be to establish a structure within the subunit which is appropriate to the task, and the nature of the intervention well depend on the specifics of the task. In this case, we have no information about the task proposed for this sub-unit, but consider the following guess, offered only for the purpose of illustration:

> "In an interdisciplinary work setting, where responsibility for the task is shared by individuals with different professional backgrounds, but where individuals have strong identifications with the team to which they are assigned, it is easy for individuals to lose their identification with the disciplines in which they

were trained. In such a case, the individuals may lose some of the skills that they acquired in their training and may certainly not keep up with advances in their field. Therefore, it is important that members of an interdisciplinary Team also be members of a discipline-oriented group, in which the focus is on maintaining one's professional identity, sharpening one's sense of the unique contribution to the overall task made by members of that discipline, and to maintaining contact with developments in other parts of the discipline."[22]

Once the task has been defined, one could begin to work with the group to establish a structure that would facilitate its performance. In this case, leadership is clearly an issue, as this is a sub-unit whose leader (M_{17}) has been rather convincingly identified as an "outcast", and the group would clearly need to identify an acceptable leadership arrangement. At the same time, one would also want to identify a set of activities in which the group members could engage in order to accomplish the task, and the nature of these activities may suggest qualifications

---

[22] This concept of the task of the Social Work/ Psychiatric Unit is essentially what Miller and Rice (1967) called a "scientific base."

for leadership of the group. In the example given here, a discipline-oriented professional activity such as a case conference would be one possible activity and this might, in turn, suggest that the group leader should be effective and stimulating in such a setting. In some such manner, the consultant would work with the group members to create a situation in which a shared understanding of the task, a system of activities for the members, and the relationships between the leaders and the members all support the accomplishment of the group's work.

Further Development of the Techniques of Intervention

The potential advantage of the similarity technique is that it allows a consultant to obtain a sharper, more detailed picture of the distribution of sentience across the range of subgroups within the organization with which it is possible for individuals to identify. This sharper picture may spur the development of techniques of intervention in two ways. First, it may focus attention on specific problems that might not have been so clearly otherwise, or characterize those problems in a new way. The identification of the "outcasts" is probably the best example of this. Though it would certainly come as no surprise to a consultant that there might be a few individuals in any given organization who are having

difficulty taking up their roles, and although a consultant who had done a number of interviews in the interdisciplinary health care team might very well have identified those individuals, the similarity technique puts the question into the foreground in a way that other approaches to organizational diagnosis might not.

New diagnostic techniques can also spur the development of intervention technique when they can be used as outcome measures- that is, when the diagnostic technique can be reapplied to measure the impact of an intervention. Again, this is an issue on which the studies reported here have nothing to say. It seems obvious that longitudinal studies in which similarity data is gathered at different points in time and in which the resulting configurations are interpreted in light of the impact of events are needed in order to explore the true potential of this technique. It seems likely that the bodies of diagnostic and intervention technique will coevolve over time.

The Problem of the "Source Representation"

The final topic to be taken up in this discussion is the question of the "source representation"- that is, the question of what, if anything, the results reported here

will permit us to conclude about the internal mental representations which serve as the basis for the similarity judgments made by our subjects. In much of what has been said above, there may appear to be an implicit assumption that the configurations that are reconstructed by the computer algorithms can be taken to be the internal representations that were used in generating the similarity ratings. In support of this assumption, one could point to their ability to account for very respectable portions (approaching 90%) of the total variance in the similarity data, suggesting that they contain almost all of the information contained in that data. Furthermore, each of the models that we have considered here (the tree models and the spatial models) appear to capture some features of the relationships between individuals that the other is either incapable of, or has difficulty in, representing. One might, therefore, conclude that a slightly better model might be found if one could combine some of the desirable features of each model into a "hybrid model" (Carroll, 1976), and that such models would be even more effective in accounting for similarity variance. If that is our criterion for deciding whether our models can be taken to be equivalent to the "source representation", then we have good reason to be satisfied.

Consider, however, the following argument, which will

be referred to, below, as the "Lyon-Venice argument."

Imagine that we need to know whether the population of Lyon, France was greater than the population of Venice, Italy, and that we decide to telephone a friend and ask him the question. Our friend then puts the phone down, goes off for a few minutes, and returns to tell us, confidently, that Lyon is the larger of the two cities. We accept the answer and hang up.

Now, suppose someone were to ask us what source our friend used to give us this answer. It would seem that this is a question that we could not answer with any certainty. For example, many maps indicate the size of cities by type of a certain size, and our friend might have looked the two cities up on a map of that sort. He might, on the other hand, have consulted an almanac or a similar reference book which contained a table of the populations of major cities. Both of these sources are capable of providing answers for the question that we asked, but they are quite different representations, developed for quite different purposes.

Our difficulty would not be alleviated by asking our friend to make a large number of such comparisons of the populations of other cities. No matter how many such questions we ask, all we will know is that the source he is

using contains that sort of information, but the resulting inquisition will not leave us any the wiser as to which of the many different sources which have that capacity he is actually using.

Now, suppose that we have asked a large number of such questions- that, in fact, we have asked him to compare the populations of every possible pairing of a given set of cities, and that we have kept a record of the responses. We could, through a process similar to those that we have considered in these studies, construct an ordering, by size of population, of the set of cities that we have used. The set could then be placed in a definite ordered sequence along a line. The representation that we had formed in this manner would explain 100% of the data that we had collected in the course of the interrogation of our friend. Would we be justified in suggesting that this representation was the one that he used?

Obviously not. In this case, it is clear that any of the other sorts of representations (maps, for example) would be as good as the representation that we had constructed. In fact, it seems clear that the goodness-of-fit that our serial ordering has achieved is due to the fact that it fits the question we have asked so well, not because it necessarily reflects the structure of the "source

representation" used by our friend. According to this argument, then, all that we can conclude is that the "source representation" must contain the information elicited by our questioning. In other words, the fact that a given representation can answer a given question tells us that it contains that information, but tells us nothing about what other information it might also contain, or how that information is organized[23]. Applied to the results reported here, the "Lyon-Venice argument" suggests that all that we can conclude is that the "source representation" contains the information that appears in the configurations that we construct from the similarity data, but that we cannot conclude that this information is the only information that is contained there or that it is particularly central information.

If we wished to pursue the question of the nature of the "source representation", we could employ one or both of the following methods. One would be to ask different sorts of questions and see which ones our friend could use his source to answer. For example, while a map would provide a basis for answering questions about the relative distances

---

[23] That this may leave out a great deal of the nature of the "source representation" is especially clear in the case of the map. Maps contain so much information that it would seem that their ability to answer "size of population" inquiries is almost incidental.

of the cities in our set, a table of populations would not. On the other hand, a table of populations would enable our friend to tell us the exact difference in population, and a map would not. In this way, we could begin to explore the full capacities and limitations of our friend's source. (In doing this, however, we would be assuming that our friend used the same source to answer all of our questions. If he had several sources available, and used different ones depending on the question that we asked, we might be fooled.)

The other method would be to bring to bear other kinds of information that we possessed about our friend and about the kinds of representations that were available to him to narrow the range of possible "source representations." If we knew that he was severely dyslexic, for example, and had great difficulty reading numerical tables, we would be more likely to guess that he had consulted a map in responding to our questions. Or we might happen to know that he had recently acquired a splendid new World Almanac and had been taking great delight in exploring it, in which case we would guess that he had used our questions as yet another opportunity to display its capacities!

Using these two methods together, we would stand a good chance of guessing our friend's source. In the studies

reported above, however, we have not used them. Instead, we have restricted ourselves to asking the same question about similarity over and over again, many times. It would seem wise, then, for us to be extremely cautious in drawing any conclusions about the nature of the "source representation" for the similarity ratings that we obtained until we have been able to ask different questions of the same "source representation" or can bring to bear other information that will narrow the range of possibilities. Then, we could begin to make guesses (formulate hypotheses) about the nature of the "source representation" and test these hypotheses against the body of data that has been accumulated, of which the studies reported here will be but a single item.

One should note, however, that nothing in this line of argument, of course, detracts from the practical use of this method in "clinical" applications. In such applications, the main point is to find out where the members of the group have made their investments of sentience. This is comparable to eliciting the full set of relative populations in the hypothetical example above, and is a question that is independent of the nature of the "source representation."

This caution is certainly the prudent course, and additional investigations are obviously desirable. Consider

the following counterargument which suggests that: 1) the "Lyon/Venice argument" overlooks an important feature of the experimental paradigm used in the studies reported here; and that 2) reconstructions of the type presented here more useful in determining the nature of the hypothetical "source representation" than the argument above would suggest.

In these studies, the counterargument would go, we did not request a specific sort of information by defining a specific type of similarity. It is as if we said to our friend "I'm going to ask you some questions about cities. Get a reference that you think will best help you answer them... Now, Lyon and Venice, are they the same or different?" In this situation, we might imagine that the criteria that he uses in answering are suggested by the representation that they were using to make the judgments, somewhat in the way that a magician forces his subject to choose the particular card that he wishes, by making it salient. If our friend has chosen a table of city populations, he is almost certain to give answers framed in those terms. Population size is not only a salient feature of such a representation, it is almost the only feature. If he has chosen a map, however, he has many possible interpretations of "same or different" to choose from, some of which (e.g. distance) will be particularly salient. In other words, with an unstructured question such as that

asked in the studies reported here, there is a greater opportunity for the nature of the representation, here thought of in terms of what kinds of information are particularly salient or particularly explicit, to exert an influence on the kinds of answers that are given. If this were to be the case, then, one might be justified in drawing inferences about the nature of the "source representation" from the sorts of reconstructions that we have considered here. The features that emerge from our reconstructions from the similarity data may not be relatively incidental parts of the "source representation", in the way that population size is a relatively incidental sort of information on a map. They may, instead, be features that are central to the character of that representation, in the way that distance is essential to the character of a spatial map.

If one believed that this was the case, then one would be inclined to build one's initial guesses (hypotheses) about the nature of the representations around these features, though it would still be wise to be cautious in concluding that these sorts of reconstructions are a sort of "royal road" to the nature of the underlying representations. Consider the following further counterargument, which raises questions that cannot be resolved by the data currently available.

Even if one accepts that, in asking the sort of unstructured similarity question that was used in these studies, the investigator has not imposed his own notion of similarity on the situation, there is still the possibility (the likelihood, even) that the subjects themselves made their own choices about what similarity meant. In so doing, they have assumed a role very much like the interrogator in the "Lyon/Venice argument", with respect to their own internal representations. If this was the case, then the (highly speculative) possibility that the representation itself can "exert an influence on the kinds of answers that are given", as suggested above, is greatly diminished, if not completely eliminated.

At this point, no conclusive answer can be given to these questions. There is no data available in these studies on what the subjects really do (much less what they do not do) as the respond to these similarity questions. All that we have is what the product of those activities. The only way to proceed, then, is with some combination of the two methods mentioned above: 1) varying the demands on the source representation, to see what sorts of questions it can answer (provided, as noted above, that one can be sure that one is making inquiries of the same "source representation" in every case); and/or 2) bringing other

information to bear on the problem. Together, these methods can be used to develop hypotheses about the nature of the source representations which can then be tested through further experimentation.

BIBLIOGRAPHY

Arabie P. & Carroll, J.D. (1980). MAPCLUS: A mathematical programming approach to fitting the ADCLUS model. Psychometrika, 45, 211-235.

Attneave, F. (1950). Dimensions of Similarity. American Journal of Psychology, 63, 516-556.

Beals, R. & Krantz, D.H. (1967). Metrics and geodesics induced by order relations. Mathematische Zeitschrift, 101, 285-298.

Beals, R., Krantz, D.H., & Tversky, A. (1968). The foundations of multidimensional scaling. Psychological Review, 75, 127-142.

Bion, W.R. (1961). Experiences in Groups. London: Tavistock.

Breiger, R.L., Boorman, S.A., & Arabie, P.A. (1975). An algorithm for clustering relational data with applications to social network analysis and comparison with multidimensional scaling. Journal of Mathematical Psychology, 12, 328-383.

Carroll, J.D. (1976). Spatial, non-spatial, and hybrid models for scaling. Psychometrika, 41, 439-463.

Carroll J.D. & Arabie, P. (1983). INDCLUS: An individual differences generalization of the ADCLUS model and the MAPCLUS algorithm. Psychometrika, 48, 157-169.

Carroll J.D. & Chang, J.J. (1970). Analysis of individual differences in multidimensional scaling via an N-way generalization of "Eckart-Young" decomposition. Psychometrika, 35, 283-319.

Corter J.E. & Tversky, A. (1987). Extended similarity trees. Psychometrika, 51, 429-451.

Horan, C.B. (1969). Multidimensional scaling: Combining observations when individuals have different perceptual structures. Psychometrika, 34, 139-165.

Johnson, S.C. (1967). Hierarchical clustering schemes. Psychometrika, 32, 241-254.

Jones, L.E. & Young, F.W. (1971. The structure of a social environment: A longitudinal individual differences scaling of the L.L. Thurstone Psychometric Laboratory (Report No. 91). Chapel Hill, N.C.: The L.L. Thurstone Psychometric Laboratory, University of North Carolina.

Jones, L.E. & Young, F.W. (1972). Structure of a social environment: Longitudinal individual differences scaling of an intact group. Journal of Personality and Social Psychology, 24, 108-121.

Katz D. & Kahn, R.L. (1966). The Social Psychology of Organizations. New York: Wiley.

Kruskal, J.B. (1964). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. Psychometrika, 29, 1-27.

Kruskal J.B. & Wish, M. (1978). Multidimensional Scaling. Beverly Hills: Sage.

McGee, V.E. (1968). Multidimensional scaling of N sets of similarity measures: A nonmetric individual differences approach. Multivariate Behavioral Research, 3, 233-248.

Marr, D. (1982). Vision. San Francisco: W.H. Freeman.

Miller E.J. & Rice, A.K. (1967). Systems of Organization. London: Tavistock.

Rice, A.K. (1958). Productivity and Social Organization. London: Tavistock.

Ross, R.T. (1934). Optimum orders for the presentation of pairs in the method of paired comparisons. Journal of Educational Psychology, 25, 375-382.

Sattath S. & Tversky, A. (1977). Additive similarity trees. Psychometrika, 42, 319-345.

Shepard, R.N. (1957). Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. Psychometrika, 22, 325-345.

Shepard, R.N. (1958a). Stimulus and response generalization: Tests of a model relating generalization to distance in psychological space. Journal of Experimental Psychology, 55, 509-523.

Shepard, R.N. (1958b). Stimulus and response generalization: Deduction of the generalization gradient from a trace model. Psychological Review, 65, 242-256.

Shepard, R.N. (1962a). The analysis of proximities: Multidimensional scaling with an unknown distance function. Part I. Psychometrika, 27, 125-140.

Shepard, R.N. (1962b). The analysis of proximities: Multidimensional scaling with an unknown distance function. Part II. Psychometrika, 27, 219-246.

Shepard, R.N. (1974). Representation of structure in similarity data: Problems and prospects. Psychometrika, 39, 373-421.

Shepard R.N. & Arabie, P. (1979). Additive clustering: Representation of similarities as combinations of discrete overlapping properties. Psychological Review, 86, 87-123.

Sokal R.R. & Sneath, P.H.A. (1963). Principles of Numerical Taxonomy. San Francisco: W.H. Freeman.

Trist E.L., Higgin, G.W., Murray, H., & Pollock, A.B. (1963). Organizational Choice. London: Tavistock.

Torgerson, W.S. (1952). Multidimensional scaling: I. Theory and method. Psychometrika, 17, 401-419.

Torgerson, W.S. (1965). Multidimensional scaling of similarity. Psychometrika, 30, 379-393.

Tucker L.R. & Messick, S.W. (1963). Individual difference model for multidimensional scaling. Psychometrika, 28, 333-367.

Tversky, A. (1977). Features of similarity. Psychological Review, 84, 327-352.

Tversky A. & Krantz, D.H. (1970). The dimensional representation and the metric structure of similarity data. Journal of Mathematical Psychology, 7, 572-597.