Cognitive Psychology **42,** 158–216 (2001)
doi:10.1006/cogp.2000.0746, available online at http://www.idealibrary.com on IDEAL®

To Matrix, Network, or Hierarchy: That Is the Question

Laura R. Novick and Sean M. Hurley

Vanderbilt University

The authors present a structural analysis of three spatial diagrams—matrices, networks, and hierarchies—that specifies 10 properties on which these diagrammatic representations are hypothesized to differ: global structure, building block, number of sets, item/link constraints, item distinguishability, link type, absence of a relation, linking relations, path, and traversal. Each property has a ''value'' for each diagram, and these property values constitute the applicability conditions for the representations. Twenty-three college students (computer science majors and math educators) selected the type of diagram they thought would be most efficient for organizing the information in each of 18 short scenarios and verbally justified the reasons for their selections. The verbal protocols were coded with respect to the structural analysis. Both the representation selection and verbal justification data provided strong support for the structural analysis. Additionally, a factor analysis of students' justifications indicated that the organization of their knowledge is consistent with the structural analysis. Students' use of the structural properties to select appropriate representations and to justify those selections indicates that the structural analysis has psychological force. © 2001 Academic Press

> Visual displays ''are essential to how scientific objects and orderly relationships are revealed and made analyzable.'' (Lynch, 1990, p. 154)

The topic of reasoning with diagrams currently is attracting interdisciplinary attention among computer scientists, philosophers, and psychologists

Portions of this research were presented at the 38th and 39th annual meetings of the Psychonomic Society (November 1996, November 1997), the 18th annual conference of the Cognitive Science Society (July 1996), and the Third International Conference on Thinking (August 1996).

We thank Melissa Francis for her helpful discussions concerning the structural analysis, for her help in creating the experimental materials, and for her help in collecting the data. We thank Doug Morse for helping code the verbal protocols and for helping us understand what Vanderbilt computer science students learn about the three spatial diagrams in their courses. We give special thanks to Lynnette Henderson, Emily Jones, and Beth O'Shea for conversations and curriculum guides that led us to the examples of matrices and networks being used by many different cultures over the past 3000 to 4000 years. J. J. Konopnicki provided thoughtful comments on an earlier version of the manuscript.

Address correspondence and reprint requests to Laura R. Novick, Department of Psychology and Human Development, Box 512 Peabody, Vanderbilt University, Nashville, TN 37203. E-mail: Laura.Novick@vanderbilt.edu.

0010-0285/01 $35.00

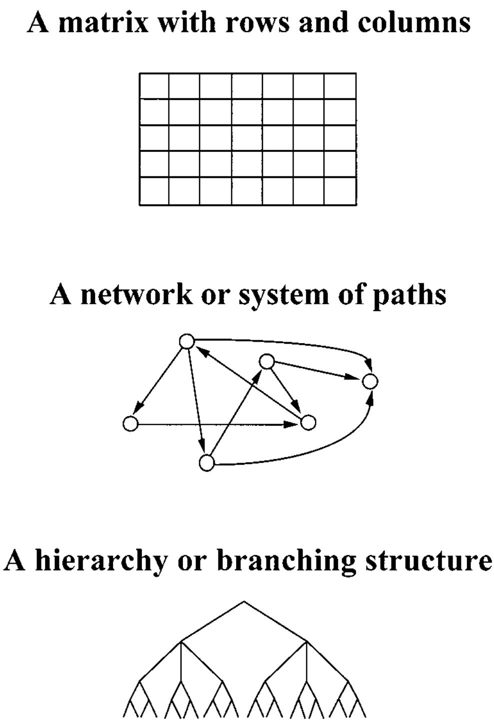

FIG. 1. The three spatial diagram representations. From ‘‘Evidence for abstract, schematic knowledge of three spatial diagram representations,’’ by L. R. Novick, S. M. Hurley, & M. Francis, 1999, *Memory & Cognition,* **27,** p. 290. Copyright 1999 by the Psychonomic Society, Inc. Adapted with permission.

(e.g., Allwein & Barwise, 1996; Glasgow, Narayanan, & Chandrasekaran, 1995). Schematic diagrams, which depict abstract concepts, are of particular interest (e.g., Butler, 1993; Hegarty, Carpenter, & Just, 1991). Examples from this category include circuit diagrams, Venn diagrams, and hierarchical trees. Schematic diagrams are important tools for thinking (e.g., Dufour-Janvier, Bednarz, & Belanger, 1987; Kindfield, 1993/1994; Larkin & Simon, 1987; Novick, Hurley, & Francis, 1999; Polya, 1957; D. Schwartz, 1993; Tversky, 2001) because they (a) simplify complex situations (Lynch, 1990; Winn, 1989), (b) make abstract concepts more concrete (Winn, 1989), and (c) substitute easier perceptual inferences for more computationally intensive search processes and sentential deductive inferences (Barwise & Etchemendy, 1991; Larkin & Simon, 1987).

This article focuses on three types of schematic diagrams—matrices, networks (path diagrams),[1] and hierarchies (see Fig. 1)—that we refer to as

[1] Technically, a network is a special type of graph in which the links between nodes are mapped into a set of real numbers (i.e., conceptually, the links are labeled with num-

spatial diagrams (e.g., Novick et al., 1999). Like other schematic diagrams, these three diagrams facilitate learning and problem solving (e.g., Bartram, 1980; Broadbent, Cooper, & Broadbent, 1978; Carroll, Thomas, & Malhotra, 1980; Day, 1988; Guri-Rozenblit, 1988; Holliday, 1976; McGuinness, 1986; Novick & Hmelo, 1994; Scanlon, 1989; S. Schwartz, 1971; Vessey & Weber, 1986).

Schematic diagrams typically rely on convention to depict both the components of the situation being represented and their organization, and the conventions must be learned before the diagrams can be understood and used successfully (Dufour-Janvier et al., 1987; Hegarty et al., 1991). The primary goal of the research presented here is to identify and investigate the ‘‘conventions’’ underlying the use of matrices, networks, and hierarchies. We propose and test an integrated structural analysis of these three diagrams. The structural analysis consists of 10 properties on which the three diagrams are hypothesized to differ. The properties are to be construed as variables for which the diagrams have different (nominal) values. These values constitute the applicability conditions for the diagrams, indicating the types of situations for which each type of diagram is best suited.

The Need for a Structural Analysis of the Three Spatial Diagrams

Although a large body of research has shown that schematic diagrams, including the three of interest here, are powerful tools for thinking, it is important to note that superior performance is only obtained when the display format and the structure of the environment are consistent (Sanfey & Hastie, 1998). It is critical, therefore, to be able to specify the situations for which each type of diagram is best suited (also see Cheng, 1996). Toward this end, several structural or functional analyses of diagrams have been proposed. However, these prior analyses are insufficient to specify the applicability conditions for matrices, networks, and hierarchies. We argue that accomplishing this goal requires a structural, rather than functional, analysis of these diagrams.

Two structural analyses of a wide variety of graphical items (scientific drawings, bar graphs, photographs, Venn diagrams, corporate logos, etc.), depicting both concrete and abstract concepts, have been proposed. In Lohse, Biolsi, Walker, and Rueter’s (1994) analysis, separate categories refer to kinds of diagrams, such as graphs, maps, icons, network charts (same as our networks), and photorealistic pictures. In Twyman’s (1979) analysis, which included text as well as diagrams, graphic stimuli are cross-classified according to methods of configuration (e.g., linear interrupted, linear branching, matrix, and nonlinear directed viewing) and modes of symboliza-

---

bers). ‘‘Graph,’’ therefore, is the more general term (Chartrand, 1985). Because psychologists usually think of line or bar graphs when they encounter the term ‘‘graph,’’ however, we have opted to use the less ambiguous (although not precisely accurate) term ‘‘network.’’

tion (e.g., verbal/numerical, pictorial, and schematic). These analyses tell us something about the structures of the three diagrams of interest here, but they do not provide the level of detail necessary to guide our characterization of their applicability conditions.

Cheng (1996) also analyzed a wide variety of diagrams from many different content domains, which led him to propose 12 functional roles for diagrams. For example, diagrams can show spatial structure and organization (F1; e.g., blueprints), can show how an object is physically assembled (F3; e.g., Novick & Morse, 2000), and can depict state spaces (F7; e.g., the periodic table and the transition state space for river crossing problems). This analysis seems to highlight the similarities among the three spatial diagrams. For example, matrices, networks, and hierarchies all can be used to depict both spatial structure and organization (F1) and state spaces (F7), and none can be (efficiently) used to show how an object is physically assembled (F3). Clearly, an analysis that highlights similarities will not provide much guidance in specifying the situations for which one type of diagram should be preferred over another.

Perhaps a different functional analysis would be more useful for distinguishing the three spatial diagrams? We suspect not. It is our contention that the primary function of these diagrams is to represent certain types of problem structures, with each type of diagram being specialized to accommodate a different set of structural features. Lohse et al. (1994) made this same point about graphical items more generally (also see Sanfey & Hastie, 1998; Winn, 1989). Thus, we view the spatial diagrams as (external) representations. For something to be a representation, it must have four components (Markman, 1999): (a) a represented world, (b) a representing world, (c) a set of rules that map elements of the represented world to elements of the representing world, and (d) a process that uses the information in the representing world. The represented world in our case is a description of a problem to be solved or of a situation in the world. The representing world contains the three spatial diagrams as abstract forms (e.g., as in Fig. 1), along with their applicability conditions. Each diagram has its own set of conventions for mapping represented elements onto representing elements. For example, for the network representation, one maps each object (or concept) in the world onto a node in the diagram and each relation between objects onto a link between corresponding nodes (Novick, 2001; Novick & Hmelo, 1994).

It is important to distinguish two types of processes that operate on the representing world. One set of processes uses the representing world and the rules for relating it to the represented world to construct a diagram to model (some aspect of) the represented world. We provide evidence for such processes in our experiment, thereby supporting our claim that the spatial diagrams are representations.

A second set of processes extracts information from the constructed diagram to reason about the represented world. A discussion of these processes is beyond the scope of this article. Matrices, networks, and hierarchies are very flexible and highly overlapping in the problem-solving processes they can support. For example, matrices can support a series of interrelated deductive inferences or application of the relative frequency definition of probability. Hierarchies also can support the latter process. The flexibility of these representations is nicely shown in studies of representational transfer, in which subjects were able to transfer a type of spatial diagram from an example problem to a test problem even though the problems required different solution procedures (Novick, 1990; Novick & Hmelo, 1994). Thus, two problems with the same goal—for example, to compute the probability of some event using the relative frequency definition of probability—may require different diagrammatic representations.

Determining the appropriate representation to use depends on assessing the degree of fit between the structure of the information to be represented and the structures of various representations. Moreover, as we have argued previously (Novick, 2001), before people can use diagrams to solve problems, they must be able to construct diagrams that accurately convey problem structure (also see McGuinness, 1986). Just as hammers, screwdrivers, and wrenches work best in specific types of situations, so do matrices, networks, and hierarchies (see Novick, 2001, for further discussion of this analogy). It is critical, therefore, to specify the problem structures for which each type of diagram is best suited. Our structural analysis goes beyond those of Lohse et al. (1994) and Twyman (1979) because we focus in depth on three interrelated types of diagrams rather than consider all different kinds of diagrams more globally.

A Structural Analysis of the Three Spatial Diagrams

Overview. Previous researchers and theoreticians have proposed a variety of structural features of matrices, networks, and hierarchies, but these have tended to be isolated reports. That is, Person 1 mentions that matrices have feature W, Person 2 notes that networks have feature X, Person 3 discusses the virtues of matrices with respect to feature Y, and so on. There has been very little discussion of the relations among the features or the representations, even in computer science textbooks on data structures (e.g., Elmasri & Navathe, 1994; Weiss, 1999). Thus, there is no comprehensive structural analysis of any of the three spatial diagrams individually, let alone of the three together. Our integrated structural analysis consists of 10 properties and their values for each diagram. As noted earlier, the property values constitute the applicability conditions for the representations. Other investigators have discussed some of the specific property values we propose, but other property values are new to our analysis. Moreover, consideration of the properties themselves as general categories is new to our analysis.

The 10 structural properties are interrelated rather than orthogonal. Nevertheless, we believe that each property focuses on a somewhat different aspect

TABLE 1
Properties Related to the General Structure of the Three Spatial Diagrams

	Global Structure (Global)
Matrix:	All the values of one variable have the values of another variable in common (i.e., the representation expresses a factorial combination of possibilities).
Network:	The representation does not have any predefined formal structure, and it does not necessarily have a unique starting or ending node.
Hierarchy:	The representation is organized into levels, beginning with a single root node (usually located at the top or right) that branches out to subsequent levels such that the identities of the nodes at one level depend on the identities of the nodes at the preceding level.
	Building Block (BBlock)
Matrix:	A cell/box denoting the intersection or combination of value *i* on one variable and value *j* on the other variable.
Network:	Two nodes and a (directional or nondirectional) link between them.
Hierarchy:	A single node that gives rise to at least two other nodes, or at least two nodes that are narrowed down to a single node, but not both (i.e., three nodes and two directional links connecting them, arranged as a ''V'' in some orientation).
	Number of Sets (NSet)
Matrix:	The rows and columns specify values along two distinct variables.
Network:	The nodes specify values along a single variable.
Hierarchy:	This representation does not naturally suggest that the nodes are arranged into a particular number or configuration of groups.
	Item/Link Constraints (ILCon)
Matrix:	Values on the same dimension (i.e., same row or same column) may not be linked.
Network:	Any node may be linked to any other node (i.e., there are no constraints).
Hierarchy:	There may not be (direct) links between nodes at the same level or between nodes in nonadjacent levels.

of representational structure. To facilitate discussion of the properties, we have grouped them into three general categories in Tables 1–3.[2] Table 1 describes four properties that seem related to the basic structure of the dia-

[2] Our analysis of the structures of the three spatial diagrams actually specifies that the analysis would be best represented in matrix form. The 10 properties would label one dimension of the matrix, and the three representations would label the other dimension. Each cell would contain the hypothesized value for the property and representation that intersect at that point in the matrix. For our own purposes, we in fact have and use such a representation. Abbreviating the property values to conserve space, the matrix spans three pages. Writing the property values in complete sentences, as we have done here, would add at least another page to the matrix. It is not clear to us that embedding a sideways-printed table that is several pages long in a journal article is the best way to facilitate comprehension. It is one thing to have a separate set of pages turned sideways and quite another to have to continually turn a journal article sideways and back again. We have opted, therefore, to present the structural analysis in three tables in this article.

TABLE 2
Properties Providing Detailed Information about the Items and Links in the Three Spatial Diagrams

	Item Distinguishability (ItemD)
Matrix:	All of the rows have identical status (i.e., are indistinguishable except by name), as do all of the columns.
Network:	All of the nodes have identical status (i.e., are indistinguishable except by name).
Hierarchy:	The nodes at a given level have identical status, but the nodes at different levels differ in status.
	Link Type (LType)
Matrix:	In general, the links between row and column values are purely associative (i.e., they are nondirectional).
Network:	The links between nodes may be associative (i.e., nondirectional), unidirectional, or bidirectional.
Hierarchy:	The links between nodes are directional such that processing flows from one end of the representation to the other.
	Absence of a Relation (AbsRel)
Matrix:	The absence of a link between a row value and a column value typically is indicated explicitly in the representation by placing a special mark (e.g., an ‘‘X’’) in the relevant cell.
Network:	The absence of a link between two nodes must be computed in all cases because there are no constraints on which nodes may be linked.
Hierarchy:	The absence of a link between two nodes generally is indicated implicitly due to the constraints on which nodes may be linked, but it must be computed for nonlinked nodes in adjacent levels.

grams: global structure, building block, number of sets, and item/link constraints. Table 2 describes three properties that provide details about the items and links in the diagrams: item distinguishability, link type, and absence of a relation. Table 3 describes three properties that are relevant for considering the potential for movement of information through the diagram: linking relations, path, and traversal. Each property has an abbreviation, indicated in Tables 1–3, that is used in labeling some of the figures of results.

Bertin (1980, 1981) noted that (appropriate) graphic presentations of data convey information about both the global organization or pattern and the individual elements or parts (e.g., the entire hierarchical tree vs. a single leaf or branch). The properties shown in Tables 1 and 3 tend to be more global (building block and linking relations are exceptions), whereas those shown in Table 2 tend to be more local.

Before we describe each of the properties in detail, it is important to make several general points about our structural analysis. First, each of the 10 properties has a value for each of the three representations. In most cases, the values for the three representations are unique, although there are a few exceptions to this. The uniqueness of the property values across representa-

TABLE 3
Properties Related to the Potential for Movement in the Three Spatial Diagrams

	Linking Relations (LRel)
This property focuses on the links going into and out of a single node.	
Matrix:	The links associated with each row or column value depict both one-to-many and many-to-one relations in the represented world, but the existence of these (many-to-many) relations must be inferred (i.e., is not directly accessible from the representation).
Network:	Any number of lines can enter and leave each node. Thus both one-to-many and many-to-one (i.e., many-to-many) relations can be represented simultaneously.
Hierarchy:	Either a single line enters and multiple lines leave each node (i.e., all depicted relations are one-to-many) or multiple lines enter and a single line leaves each node (i.e., all depicted relations are many-to-one), but not both.
	Existence of Paths (Path)
Matrix:	This representation does not show paths connecting subsets of (more than two) items.
Network:	This representation shows paths connecting subsets of (more than two) nodes.
Hierarchy:	This representation shows paths connecting subsets of (more than two) nodes.
	Traversing the Representation (TravR)
Matrix:	It does not really make sense to talk about traversing this type of representation.
Network:	Multiple paths from one node to another are possible because closed loops are allowed in the representation.
Hierarchy:	For any pair of nodes, A and B, there is only one path to get from one to the other (i.e., closed loops are not allowed).

tions supports our claim that matrices, networks, and hierarchies are distinct representations that are optimized for different problem structures.

Second, the properties likely vary in their importance, or diagnosticity, for specifying a particular representation. For example, we suspect that the number of sets property is more important than the traversal property for specifying a matrix, whereas the reverse is the case for the hierarchy. We note in the paragraphs below where we have specific hypotheses about differences in the diagnosticity of the properties across representations. However, further consideration and empirical investigation of this issue are beyond the scope of this article.

Third, our goal is to specify the structural features for which each type of representation is best suited. Thus, our matrix features specify a prototypical matrix and similarly for the other two representations. Clearly, hybrid representations can be constructed. Moreover, any information that can be represented using one of the three spatial diagrams can be represented using either of the other two diagrams. That it is possible to do this, however, does not mean that it is desirable. For example, Chartrand (1985) noted both that the information represented in a network can easily be represented in a matrix and that there are several disadvantages to doing so. Our claim is that the three spatial diagrams are optimized to accommodate different structural features. These features are specified in our structural analysis. It will be an important topic for future research to examine problem solvers' behavior when they encounter situations whose structural features overlap two or more representations. However, fruitful investigation of this issue is predicated on having a firm grasp of the prototypes. Hence the present article addresses the logically prior issue.

Global structure. Most abstract representations have a general form, what we refer to as a global structure, that allows one to quickly recognize the type of representation illustrated without going through a detailed analysis of each of its parts. Thus, for example, we can easily distinguish between a bar graph and a line graph because they have different general forms. The same is true for the matrix and hierarchy representations (see Table 1).

As noted generally by Inhelder and Piaget (1964) and with respect to tables of data by Bertin (1980), the matrix representation has a structure in which all of the values of one variable have the values of another variable in common. That is, the representation expresses a factorial combination of possibilities. Day (1988, p. 282) stated more generally that matrices are ‘‘ideally suited to display the union of factors.’’ Tversky (2001) noted that tables cross-classify items. Factorial combination, union, and cross-classification are descriptions of the global structure of the representing world.

In contrast, for the hierarchy representation, the key terms in describing the global structure of the representing world are levels and dependency. That is, the hierarchy is organized into levels, beginning with a single root node (usually located at the top or right) that branches out to subsequent levels such that the identities of the nodes at a given level depend on the identities of the nodes at the preceding level. Bertin (1981) mentioned that a hierarchy is organized into levels, but he did not discuss the notion of dependency (also see Elmasri & Navathe, 1994; Epp, 1995). Vessey and Weber (1986) suggested the idea of dependency but did not mention the organization into levels.

What about the network representation? Unlike most types of abstract diagrams, this one is characterized by the lack of any formal structure. The network does not have a predefined structural form or even a unique starting or ending node. The nodes and links can appear in any configuration, and, as is shown with the item/link constraints property, the nodes can be connected in any way. Mathematically, the network representation is the general form (e.g., Chartrand, 1985), of which both matrices and hierarchies are special cases, so it makes sense that the network does not have a predefined formal structure. Bertin (1981) indirectly suggested that networks lack a predefined structure by illustrating a variety of different structures they can take.

Building block. Although we can recognize the different representations based on their global structures, to construct them it is necessary to know what their basic units, or building blocks, are (see Table 1). For the matrix, the building block is a cell (or box) denoting the intersection or combination of value *i* on one variable and value *j* on the other variable (Inhelder & Piaget, 1964). Put together a number of such units as specified by the global structure property and you get a matrix. For the network, the building block is two nodes and some type of link between them. The hierarchy building block can be described in either of two equivalent ways, depending on whether construction of the representation begins or ends at the root node. If construction begins at the root node, then the building block consists of a single node that gives rise to (at least) two other nodes (Elmasri & Navathe, 1994; Weiss, 1999). But if construction ends at the root node, then the building block consists of (at least) two nodes that are narrowed down to a single node. These two statements describe the same basic unit, which is three nodes and two directional links connecting them, arranged as a ''V'' in some orientation.

Number of sets. The number of sets property clearly distinguishes the matrix and network representations (see Table 1). Whereas the row and column structure of a matrix naturally suggests values along two distinct variables (Bertin, 1981; Cormen, Leiserson, & Rivest, 1990; Day, 1988), the collection of nodes in a network naturally suggests values along only a single variable (Bertin, 1981). In the represented world, this corresponds to two sets of objects or concepts versus one. This is not to say that it is impossible to use a matrix when there is only one set of objects or a network when there is more than one set of objects. In the former case, the rows and columns of the matrix could be labeled identically (Bertin, 1981). In the latter case, one can simply create a network node for each item in each set. Our claim is simply that the structure of the matrix is optimized for two sets, whereas the structure of the network is optimized for one.

In using a network to represent more than one set of objects, one loses easy access to the fact that the objects belong to different sets (also see the item distinguishability property). Using a matrix to represent only a single set of objects wastes space because the top and bottom (triangular) halves of the matrix provide redundant information. A matrix also can accommodate more than two sets of objects. For routine problem solving on paper, however (as opposed to computer programming), one must resort to creating multiple matrices, each of which represents two sets of objects. Thus, for example, three matrices would be needed to represent the relations among values along three dimensions: One matrix would represent combinations of items in sets A and B, a second matrix would represent combinations of items in sets A and C, and a third matrix would represent combinations of items in sets B and C. Each matrix, then, would fit the prototype we have specified.

Although we hypothesize that the number of sets property is important for distinguishing when to use a matrix versus a network representation, we suggest that it is not particularly important for specifying the appropriateness of a hierarchy representation. The structure of the hierarchy does not naturally suggest that the nodes are arranged into a particular number or configuration of groups. In using the mapping rules to apply a hierarchy representation to particular situations, it becomes clear that both the breakdown of items into sets and the relation of those sets to the structure of the hierarchy are specified by the semantics of the represented world. For example, if one uses a hierarchy to represent the division of land animals into mammals and reptiles, it makes sense to think of the left and right ‘‘sides’’ of the hierarchy as specifying different sets of objects (namely mammals vs. reptiles). In contrast, in representing the officers in different units on a military base, each level in the hierarchy specifies a different set, namely a different rank such as general or lieutenant. Finally, in using a hierarchy to represent the breakdown of the number ‘‘48’’ into its multiplicative factors, one would either say that there is a single set of factors or that the concept of sets does not apply.

Item/link constraints. Because the three spatial diagrams differ in their global structures, they place different constraints on the possible links between row and column values (matrix) or nodes (network and hierarchy; see Table 1). As we alluded to in our discussion of the global structure property, the network representation places no constraints on the links between nodes. Any node may be linked to any other node. This is not to say that every node must be linked to every other node. Clearly that is not the case. However, any constraints on these links that exist come solely from the semantics of the represented world.

In contrast, the matrix and hierarchy representations both place some constraints on which ‘‘items’’ are or are not allowed to be linked. In a matrix, values on the same dimension (i.e., same row or same column) may not be linked. In the represented world, this means that items in the same set may not be linked. In a hierarchy, there may not be (direct) links between nodes at the same level or between nodes in nonadjacent levels. We were not able to find any references to the item/link constraints property in the literature.

Item distinguishability. The item distinguishability property indicates whether the items in the representation differ in name only or in some other more structured way (see Table 2). For the network representation, all the nodes have identical status. That is, they are indistinguishable except by name. The same is true for all of the rows in a matrix and likewise for all of the columns. This is not to say that some order cannot be imposed on the rows and columns of a matrix or on the nodes of a network. Clearly, one can do that. For example, to best organize the information in a particular represented world, the labels for the rows or columns of a matrix could be alphabetized or ordered along some dimension, and the labeled nodes in a network could be arranged so that there is generally a chronological ordering from left to right. But these orders are suggested by the semantics of the represented world and are not inherent to the representing worlds.

Bertin (1981) seems to agree, stating that in the basic (or default) case the nodes in a network, and the rows and columns of a matrix, are ‘‘equal’’ or ‘‘reorderable.’’ In situations for which this is not true, the representation may acquire a special name. For example, when the nodes of a network are ordered according to geographical constraints, Bertin refers to the resulting representation as a topography. Elmasri and Navathe (1994) indicate that there is no preference for one ordering of the tuples (rows) or attributes (columns) of a relation (matrix) over another.

In contrast, the hierarchy representation does inherently impose an ordering on the nodes. In particular, although the nodes at a given level in the representation have identical status, the nodes at different levels differ in status. The idea of differing status is conveyed by the representing world, although the exact definition of status depends on the particular content from the represented world that is mapped onto the hierarchy form. For example, in describing a corporate power structure, status refers to ranks in an ordered chain of command. In describing the classification hierarchy for vehicles or fruit, for example, status refers to levels of inclusiveness (e.g., Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). And in describing the results of certain tournaments (e.g., the NCAA basketball tournament), status refers to the number of games won in a certain time period.

Link type. The three representations also differ in the types of links between items that they most easily accommodate (see Table 2). In the matrix representation, the links generally are purely associative (i.e., nondirectional; Elmasri & Navathe, 1994). Bertin (1980) conveyed this idea indirectly, in talking about tables of data, through the exhaustive list of possible relations he gave, none of which was directional. The exact nature of the associative link depends on the semantics of the represented world (e.g., see Bertin, 1980; Novick et al., 1999).

In contrast, in the hierarchy representation the links between nodes are directional such that processing flows from one end of the diagram to the other. Whether the links go in the direction from the root of the tree to the leaves or vice versa depends on the specifics of the represented world. D. Schwartz (1993) considers the directional nature of the links in a hierarchy to be implicit rather than explicit. That is, the interpretation of directionality depends on our knowledge about this type of representation.

Because the network is a general form that subsumes the matrix and the hierarchy (e.g., Chartrand, 1985), it is more flexible in the types of links it allows. In a network, the links may be associative (denoted by lines connecting nodes), unidirectional (denoted by single-headed arrows), or bidirectional (denoted by double-headed arrows).[3] D. Schwartz (1993) noted that

[3] In formal mathematical graphs, bidirectional links are typically indicated by two separate directional arrows rather than by a single two-headed arrow (Chartrand, 1985).

networks indicate direction explicitly (also see Cormen et al., 1990; Epp, 1995; Weiss, 1999). Epp (1995) gave an example of an undirected network in which the links denoted associative relations.

Absence of a relation. Just as they differ in the types of links between items that they most easily accommodate, the three spatial diagrams differ with respect to their ability to convey information about the absence of a link between items (see Table 2). The matrix representation is best suited to conveying such information: The absence of a link between a row value and a column value can be indicated explicitly by placing a special mark (e.g., an ‘‘X’’) in the relevant cell (Bertin, 1980, 1981; Elmasri & Navathe, 1994). Tversky (2001) and Epp (1995) talked about tables or matrices representing the absence of a relation in the context of specific examples: For example, in a train schedule a blank space rather than a time indicates that a particular train does not stop at a particular station.

The network representation, on the other hand, does not explicitly convey information about the absence of a link between two items. Rather, this information must be computed by finding the nodes in question and tracing all links out of one node to determine that none of them end at the other node. No predictions can be made *a priori* about which nodes will not be linked because, as noted earlier, there are no constraints on the linking of nodes (see the item/link constraints property).

The hierarchy representation presents an intermediate case with respect to this property. The absence of a link is never indicated explicitly (Weiss, 1999), as it is in the matrix. However, it is indicated implicitly due to the constraints on which nodes may be linked. In particular, by definition there are no links between nodes at the same level or between nodes in nonadjacent levels (see the item/link constraints property). These nonlinks are not explicitly marked, but they are predictable *a priori*. However, the absence of a link between some of the nodes in adjacent levels must be computed, as described for the network.

Linking relations. The linking relations property focuses on the links going into and out of a single item. It specifies whether the representation allows many links both entering and leaving an item (see Table 3). This property is especially important for distinguishing the hierarchy and network representations. For the network, any number of lines (i.e., relations in the represented world) can enter and leave each node. Thus, the network can depict simultaneously both one-to-many and many-to-one (i.e., many-to-many) relations in the represented world (Elmasri & Navathe, 1994; D. Schwartz, 1993).

In contrast, for the hierarchy, only one of the following two statements can be true: (a) A single line enters and multiple lines leave each node (i.e., all depicted relations are one-to-many; Elmasri & Navathe, 1994; D. Schwartz, 1993; Weiss, 1999) or (b) multiple lines enter and a single line leaves each node (i.e., all depicted relations are many-to-one). Whether the relations depicted in a hierarchy are conceptualized as one-to-many or many-to-one depends on whether one traverses the hierarchy from the root to the leaves or from the leaves to the root, respectively.

Like the network representation, the links associated with each row or column value in a matrix depict both one-to-many and many-to-one relations in the represented world. However, this does not seem to be a particularly salient aspect of this representation. In fact, the existence of these many-to-many relations must be inferred, rather than being directly accessible from the representation as for the network and hierarchy.

Path. In contrast to the linking relations property, the path property focuses on the representation as a whole. It specifies whether the representation shows paths connecting subsets of (more than two) items (see Table 3). Paths are evident in the hierarchy (Chartrand, 1985; Epp, 1995; Vessey & Weber, 1986; Weiss, 1999) and the network (Chartrand, 1985; Cormen et al., 1990; Epp, 1995; Weiss, 1999) but not in the matrix. In fact, networks are commonly referred to as path diagrams outside the mathematical literature.[4]

Traversal. Traversal, the final property, is most relevant for distinguishing the network and hierarchy representations. It specifies structural differences in the types of paths that are possible (see Table 3). For the network, multiple paths from one node to another are possible because closed loops are allowed in the representation. In contrast, for the hierarchy, there is only one path between any pair of nodes. That is, closed loops are not allowed. In fact, hierarchies are sometimes referred to as acyclic graphs (e.g., Cormen et al., 1990). Because the matrix representation does not indicate paths connecting subsets of items (see the path property), it does not really make sense to talk about traversing this type of representation.

Weiss (1999) and Epp (1995) noted that networks may contain cycles, and both gave an example showing how two specific nodes in a certain network were connected by two different paths. Bertin (1981) stated that when there is only one path between nodes the network becomes a tree (i.e., a hierarchy), thereby implying also that in the general case networks allow multiple paths between pairs of nodes. Chartrand (1985) noted that any connected graph (network) that does not have cycles (i.e., closed loops) is called a tree (also see Epp, 1995), thereby implying that networks do allow cycles. He also noted that there is only one path between any two nodes in a tree, as did Cormen et al. (1990). Broadbent et al. (1978) noted that one limitation of hierarchies is that there is only a single retrieval path for each item in the represented world.

Overview of the Present Research

As we have discussed, other theoreticians agree with a number of the property values we propose in our structural analysis of matrices, networks,

[4] Technically, a path in a network is a connected set of nodes and associated links in which no node is repeated (Chartrand, 1985). We use the term more colloquially (i.e., our definition allows nodes to be repeated), in keeping with the colloquial designation of this representation as a path diagram.

and hierarchies. To our knowledge, however, there is no empirical support for either the properties themselves or the specific values on the properties that constitute the applicability conditions for the three spatial diagrams. Therefore, in the remainder of this article, we report the results of a study that was designed to assess the validity of our structural analysis of these representations. In addition to gathering evidence for the properties individually, we investigated how the property knowledge is organized in memory. In our structural analysis, subsets of the properties are hypothesized to be closely related. For example, linking relations is the local property that enables the global differences in traversal. We might expect students to know about at least some of these relations. This question addresses the issue of whether the property knowledge in students' schemas is organized in a way that reflects the structural dependencies among the properties or whether the known properties are simply listed in some arbitrary order (or perhaps are organized in some other manner).

To assess the adequacy of our structural analysis, it is important to select as subjects people who would be expected, *a priori*, to be relatively knowledgeable about the three spatial diagrams. Such subjects give us the best chance to find support for, or contradictions to, our structural analysis as well as to discover additional distinguishing properties. We sampled college students from two such populations: advanced computer science and math education majors.

The subjects read 18 short scenarios and had to choose which of two types of spatial diagrams would be best for representing the information in each scenario. Dufour-Janvier et al. (1987) have stressed the importance of being able to select the most appropriate representation for a given task (also see Novick, 2001). Each scenario focused on the property value for one representation. For example, one scenario focused on the global structure property value for the matrix. Because our structural analysis makes a strong prediction as to which type of diagram is most appropriate for each scenario, subjects' representation choices provide one source of evidence concerning the validity of our structural analysis. In addition, for each scenario subjects were asked to verbally justify why they chose the representation they did and why they did not choose the other representation. These verbal protocols were coded for statements reflecting the hypothesized applicability conditions for the representations. These data provide more detailed information concerning the structural properties and how they are organized in memory.

METHOD

Subjects

The subjects were 23 Vanderbilt University students who were recruited from two distinct populations: math educators and computer science majors. These students were selected because they were expected, *a priori*, to be relatively knowledgeable about the three spatial diagrams due to their college coursework. They were paid for their participation.

The math educators were 11 (4 male and 7 female) of the 16 students who were student teaching in secondary-level math classes during the spring semester of 1 academic year. Eight of these subjects were seniors with a double major in secondary education and mathematics. The other three were masters-level students who were getting their certification in math education. The computer-science majors were recruited from a junior-level computer science course that had as its prerequisite three courses on data structures and related topics. In these courses, the students learned about matrices (arrays), networks (graphs), and hierarchies (trees), among many other topics. Approximately 30 of the 50–60 students in the class expressed an interest in participating in the study. The first 12 students (11 male and 1 female) we could schedule participated (there was only 1 female on our list).

Property-Focused Scenarios

Overview. The 18 scenarios were all set in a vaguely medical context. We used the same general context for all scenarios so that the cover story of a scenario would not provide a strong cue as to the best representation. We expected that if we used familiar contexts (e.g., a military power structure), subjects might base their responses on the cover stories rather than on the structural information presented (in this case choosing the hierarchy representation).

Each scenario was written to focus as purely as possible on the value for one property for one representation, e.g., the matrix value for the global structure property. The scenarios were three to five sentences (five to eight lines) long. The first sentence or two set up the story context or setting for the scenario. The next sentence or two implemented the property focusing within the designated cover story. The last sentence indicated that somebody in the story wanted a diagram for some purpose relevant to the cover story. Table 4 shows three scenarios that focus on the global structure property, with each scenario focusing on the value for a different representation.

One line below each scenario, the following sentence was printed: ‘‘Based on the information given, which of the following two types of representations do you think would be the best (i.e., most efficient) for this scenario?’’ Subjects were then given a choice between two representations, the one whose property value was focused in the scenario and a contrasting representation that had a different value on the focused property. The two choices were presented as labeled diagrams (see Fig. 1). By giving subjects a choice of only two representations, we were able to ensure that our structural analysis would make a strong prediction as to the most appropriate representation for each scenario. Each pair of representations was used equally often.

Properties focused. We selected a subset of 6 of the 10 properties to use for the scenarios. One criterion for choosing a property was that it had ‘‘positive’’ values for at least two of the three representations so that we could write at least two scenarios focusing on that property. It was difficult to write scenarios to focus on the ‘‘negative’’ property values, such as matrices are not good for showing paths among subsets of items. To increase the breadth of the property values used, a second criterion was that at least two of the positive values for a selected property had to be different. The building block property was excluded because it turned out to be quite difficult to write scenarios that focused purely on that property without very strongly suggesting either or both of the global structure and link type properties. This circumstance is consistent with our claim that the properties describe overlapping rather than orthogonal aspects of representational structure. Finally, we excluded the item distinguishability property because it seemed less important for discriminating among the three representations.

The six selected properties—global structure, number of sets, item/link constraints, link type, linking relations, and traversal—were used in 2–5 scenarios each. Across the 18 scenarios, each representation had one of its property values focused six times. Table 5 summarizes the structural characteristics of the scenarios. The entries in the table indicate the representation whose property value was focused in a scenario, the contrasting representation provided as a choice, and the (arbitrarily assigned) scenario number. The global structure property was used most often because it was considered most important. For the hierarchy representation, one

TABLE 4
Three Property-Focused Scenarios for the Global Structure Property

Matrix: value = factorial combination of possibilities

(Scenario 10) In the psychiatric ward of a certain hospital, each patient sees only one doctor, who is responsible for diagnosis and treatment. A researcher is interested in determining whether patients would receive different diagnoses from different doctors. Therefore, she selected a group of newly-admitted patients and asked all of the staff psychiatrists to submit a diagnosis for each patient in the group. The department chair would like a diagram showing each doctor's diagnosis for each patient.

Network: value = no formal structure

(Scenario 2) Sisters of Mercy Hospital serves a close-knit religious community that observes strict laws forbidding the exchange of blood with people who are not members of their sect. These complicated laws also govern the exchange of blood among community members (e.g., unmarried women over the age of 20 cannot donate blood to married women). Although the constraints on who can donate blood to whom are well-specified by law, they do not follow any coherent pattern. The hospital would like a diagram showing who in the community may donate blood to whom.

Hierarchy: value = organization into levels

(Scenario 15) In her next lecture, a pharmacy professor plans to discuss the composition of a new over-the-counter flu medicine. The medicine contains several active ingredients and several inactive ingredients. Each of these ingredients can be broken down into its basic chemical structure. Thus the composition of the flu medicine can be described at a variety of different levels. The pharmacy professor would like a diagram showing the composition of this medicine.

Note. The sentence that implements the property focusing is underlined in the table, although it was not underlined for the subjects.

TABLE 5
Focused Properties and Contrasting Representations for the 18 Property-Focused Scenarios

	Representation value focused on		
Property	Matrix no.	Network no.	Hierarchy no.
Global Structure	10 (H contrast) 5 (N contrast)	2 (H contrast)	12 (M contrast) 15 (M contrast)
Number of Sets	8 (N contrast)	3 (M contrast)	
Item/Link Constraints	13 (N contrast)	9 (M contrast)	7 (N contrast)
Link Type	6 (H contrast)	18 (M contrast)	
Linking Relations	17 (H contrast)	11 (H contrast)	16 (N contrast)
Traversal		4 (H contrast)	1 (M contrast) 14 (N contrast)

Note. M, matrix; N, network; and H, hierarchy.

TABLE 6
Three Additional Property-Focused Scenarios

Matrix: Number of Sets

(Scenario 8) In treating a certain illness caused by a malfunctioning pituitary gland, endocrinologists typically choose two drugs from among those that belong to Category A or Category B. These two types of drugs have very different properties, and it is important for the proper treatment of the illness to maintain the distinction between the two categories of drugs. The American Endocrinology Association would like a diagram showing which drugs can be prescribed together.

Network: Linking Relations

(Scenario 11) Because emergency medical technicians must keep sharp on several important skills that they rarely need to use, the director of the emergency unit in a small city has implemented a new evaluation program. To ensure that the testing is conducted fairly, two technicians cannot evaluate each other. According to the rules the director set up, each technician must be evaluated by at least two other technicians. In turn, he or she may evaluate any number of other technicians. The director would like a diagram showing who evaluated whom during the last testing period.

Hierarchy: Traversal

(Scenario 14) Researchers are trying to create more effective drugs to treat a certain form of cancer. To do this, they start with the most promising drug to date, Xantalin, and modify its chemical structure slightly in a variety of ways to create several new drugs. These drugs are then further modified in a variety of ways to create yet more potential cancer drugs, etc. It turns out that if you trace each of the new drugs back to Xantalin, you will find that there is only a single sequence of modifications through which each can be synthesized. The head of the oncology research division would like a diagram showing how all of the new drugs can be synthesized from Xantalin.

Note. The focused-on property in each scenario is underlined.

of the global structure scenarios focused on the idea of levels (no. 15) and the other focused on the idea of dependency (no. 12). Three of the global structure scenarios are shown in Table 4. Three additional scenarios, that focus on three different properties, are shown in Table 6.

Counterbalancing information. The order in which the two representations were printed on the response page was counterbalanced so that each representation appeared equally often on the left and right sides. Each representation was the focused one approximately equally often in each position. Overall, the left and right positions contained the predicted best representation 8 and 10 times, respectively. It was not possible to fully counterbalance everything, and we deemed it most important to have each representation appear equally often on each side when paired with each of the other representations.

The 18 scenarios were collated into booklets in two different orders, which were random except for several constraints: (a) The same property could not be focused on for two scenarios in a row, (b) the same pair of representations could not be used for more than two scenarios in a row, (c) a particular representation could not appear in the same position (left or right) for more than three scenarios in a row, (d) a particular representation could not be the predicted correct answer for more than two scenarios in a row, and (e) a particular representation could not be used for more than four scenarios in a row (regardless of the representation with which it was paired). These constraints were intended to minimize the extent to which subjects' responses to a scenario might be guided by superficial aspects of the immediately preceding scenario.

Design

Although we selected subjects from two populations—computer science majors and math educators—we had no basis for predicting specific differences between them. For this reason, and because group differences were not the focus of this study, all 23 subjects were considered as a single group. The main factor of interest was scenario, which was manipulated within subjects. It is possible to provide alternative conceptions of this factor, such as the property being focused, the predicted best representation, or the pair of representations being contrasted. We consider different conceptions of this factor in our data analyses, depending on the issue being examined.

Procedure

Subjects participated individually. They first completed a task designed to help them access their general knowledge about the three spatial diagrams. For each representation, subjects were given 3 min to generate as many examples as they could of situations for which that type of representation would be useful for organizing information. Novick et al. (1999) found that this is an effective cue to students' abstract knowledge about these representations. Across subjects, all six possible orders for presenting the three representations were used. For each subject group, each presentation order for this task was combined with each of the two scenario booklet orders (except that one of the presentation orders was used with only one scenario booklet order for the math educators because there were only 11 subjects in that group).

After completing the example generation task, subjects were given the booklet with the 18 property-focused scenarios and a tape recorder was started. The procedure for each scenario was as follows. First, subjects read the scenario out loud. Then they picked which of the two displayed types of representations they thought best captured the structure of the scenario and rated their confidence in their selection. Confidence was rated on a 4-point scale with 1 meaning *not at all sure* and 4 meaning *very sure.* Then subjects justified their representation choice out loud. To help motivate subjects to fully explain their choices, they were asked to frame their reasons as if they were talking to a junior high or high school student, say a ninth-grader. They were told to think of their task as one of trying to help this hypothetical student understand the kinds of situations for which each type of representation would be most appropriate. After subjects had given their justification, they were asked the following question: ''OK, now what can you tell the junior high or high school student about why you think the [name of the unselected representation] would be a less good, or less efficient, type of representation for this situation?'' After the first several scenarios, many subjects answered this question in their initial justification response and therefore did not need to be prompted. Other subjects waited until the question was asked for every scenario before providing this information.

Throughout the experiment, subjects sat facing the experimenter to facilitate communication between them. While the subjects were giving their verbal protocols, the experimenter was very careful not to give (verbal or nonverbal) feedback on their responses or to answer any questions except procedural ones. Subjects sometimes asked whether they were interpreting a particular sentence in a scenario correctly. In these situations, the experimenter indicated that he/she could not provide any additional information, and the subject just had to do the best he/she could with the information provided. A recurring procedural question was whether subjects were allowed to change their representation choice after they had started giving their justification. The experimenter always answered ''yes'' to this question.

RESULTS

This study was designed to accomplish two goals. First, and most importantly, we wanted to evaluate the adequacy of our structural analysis of matrices, networks, and hierarchies. Second, we wanted to learn about how the structural properties are organized in students' knowledge bases and which properties are perceived to be most important for each type of representation. These topics will be addressed in separate subsections. We begin with a discussion of how the protocols were coded because those data inform all of the issues.

Coding the Verbal Protocol Data

In this section, we provide an overview of the protocol coding scheme and the method of coding the protocols. We also present data on the reliability of the coding scheme. The structural codes are described in detail in Appendix A. Additional information about how the coding scheme was developed and how the protocols were coded may be found in Appendix B. Some examples of coded protocols are provided in Appendix C.

Coding criteria were developed for each of the values of the 10 structural properties. In most cases, the coding criteria were simple elaborations of the property values given in Tables 1–3. We found that the item distinguishability property was rarely coded. Therefore, it was removed from the coding scheme and the affected statements were recoded before computing the reliabilities. It is unclear whether (a) the item distinguishability property is not important for distinguishing among the three representations, contrary to our structural analysis; (b) our subjects did not know about this property; or (c) our scenarios did not present situations for which subjects deemed this property important for choosing the most appropriate type of representation. It is not possible to distinguish among these alternatives with the present data.

Six additional coding categories were created to capture nonstructural aspects of subjects' justifications (e.g., self-monitoring statements). For our purposes here, the most important of these codes was one indicating that subjects used one of the presented diagrams to represent the information in the scenario. That is, they described associating specific terms from the scenario with specific parts of a matrix, network, or hierarchy. Finally, unintelligible statements and statements that did not fit into any of the other categories were placed in a ''garbage'' category.

Each subject's justification for a particular scenario received one or more different codes ($M = 4.83$, $SD = 1.23$). The codes were given based on evidence from a coherent group of statements, which we referred to as a bracketed unit, but they applied to the scenario as a whole. Therefore, although two or more brackets for a particular scenario might have received identical codes, each code was counted only once for that scenario for analytical purposes (but it could be counted again for a different scenario).

Reliabilities were based on a random sample of 13 of the 23 subjects (7 from computer science and 6 from math education). They were computed separately for each coding category (e.g., each structural property), but collapsed across the values for that category and the representation to which the code referred. The two coders were said to agree if they gave exactly the same code (i.e., the same representation and property value) for a particular subject and scenario. The reliabilities were computed as the proportion of times the two coders agreed that a particular code should be given. For each category, we divided the number of codes agreed upon by the total number of codes given in that category. For example, for the 13 reliability subjects, the two coders combined gave 218 global structure codes. Both coders used the same global structure code for 88 scenario justifications, yielding 176 such codes participating in agreements. For 42 scenarios, only one of the two coders gave a global structure code, yielding 42 such codes participating in disagreements. The proportion agreement, therefore, is 176/218 = .81.

This method of computing reliabilities is very conservative because scenarios for which both coders agreed that a particular coding category was *not* warranted based on what subjects said are not counted in the computation of reliabilities. Inclusion of these agreements would greatly increase the estimates of the reliabilities of the coding categories. For example, consider again the global structure category. The reliability analysis was based on the coding of 18 scenarios for each of 13 subjects for a total of 234 scenario justifications. A particular coding category could be used twice for each justification, once for the representation chosen by the subject and once for the representation not chosen. Thus there were 468 opportunities for the two coders to agree or disagree.[5] Using the opportunities as the unit of analysis now, there were 42 disagreements, 88 ''positive'' agreements, and 338 ''negative'' agreements, yielding a reliability of 426/468 = .91 (compared to a reliability of .81 based on just the codes given).

Table 7 shows the number of times each coding category was used by both coders combined (for the reliability sample of 13 subjects) and the proportion of times the two coders agreed (i.e., the reliabilities). Overall, the two coders agreed on 74% of the structural codes given, which is quite good considering the number of codes and the complexity of the coding system. Reliabilities for the individual structural categories are adequate for all categories except building block, which we discuss in Appendix B. Note, however, that if the alternative method of computing reliabilities is used, in which negative agreements are counted, the reliability of this category increases from .62 to .89. Reliabilities for the nonstructural categories are quite good, except for the ''garbage'' category, which seems reasonable. In any case, no conclusions are drawn about that category. To resolve the coding discrepancies, Coder 1 recoded the bracketed units that led to disagreements.

[5] Actually, a particular category could be used more than once for a representation in certain cases, such as when negated codes were allowed. For example, a subject's justification for a scenario could receive the following three codes: H/Global/H, M/Global/∼H, and M/Global/M. Thus our statement that there were 468 opportunities for the coders to agree or not is a slight underestimate. In practice, this underestimate makes this measure of reliability slightly more conservative than it would otherwise be.

TABLE 7
The Total Number of Codes Given by Both Coders Combined and the Proportion Agreement (Reliability; Conservative Method) for Each Coding Category (for the Reliability Sample of 13 Subjects)

Coding category	Total codes (both coders)	Proportion agreement
Global Structure	218	.81
Building Block	138	.62
Number of Sets	83	.84
Item/Link Constraints	24	.75
Link Type	267	.71
Absence of a Relation	41	.93
Linking Relations	66	.79
Path	87	.67
Traversal	46	.70
All structural codes	970	.74
Use Diagram to Represent Scenario	459	.94
Five Other Nonstructural Codes	556	.84
''Garbage''	68	.47
All codes	2053	.80

Subjects Use the Spatial Diagrams as Representations

Recall that the definition of a representation (Markman, 1999) has four components: a represented world, a representing world, a set of rules for relating the two, and processes that use all of the above. Without the processes, there is only the possibility of representation and not actual representation. Our protocols provide clear evidence for representational processes. Our code for using the diagrams as representations was the most frequent code given. Across the 18 scenarios, subjects averaged nearly 20 diagram use codes ($M = 19.91$, $SD = 5.91$). Subjects typically described how the diagram they selected as most appropriate for a scenario would be used to represent the information in that scenario ($M = 14.30$ of 18 scenarios across subjects), and they sometimes described how the other diagram would be used when explaining why it was not as good a choice ($M = 5.61$ of 18 scenarios). Each of the spatial diagrams had one of its property values focused on in six scenarios. The average number of diagram use codes for each type of diagram was 8.17 for the matrix, 5.70 for the network, and 6.04 for the hierarchy.

Evaluating the Structural Analysis of Matrices, Networks, and Hierarchies

There are several ways to evaluate the structural analysis of the three spatial diagrams presented in Tables 1–3. One method involves examining subjects' representation choices. Because each scenario focused on the unique

TABLE 8
Mean Levels of Conformity to the Predictions for Each Representation and Each Structural Property and Tests against Chance Responding

Category	No. of scenarios	Proportion choosing predicted spatial diagram	Test against chance (=.50)
Representation			
Matrix	6	.96	$t(22) = 25.74$, $MS_e = .0180$
Network	6	.70	$t(22) = 5.85$, $MS_e = .0347$
Hierarchy	6	.96	$t(22) = 31.64$, $MS_e = .0147$
Property			
Global Structure	5	.94	$t(22) = 18.84$, $MS_e = .0233$
Number of Sets	2	.80	$t(22) = 5.85$, $MS_e = .0520$
Item/Link Constraints	3	.68	$t(22) = 4.08$, $MS_e = .0443$
Link Type	2	.93	$t(22) = 12.11$, $MS_e = .0359$
Linking Relations	3	.94	$t(22) = 16.41$, $MS_e = .0269$
Traversal	3	.91	$t(22) = 13.24$, $MS_e = .0312$

Note. All $p < .001$.

property value for a single representation, the structural analysis makes a strong prediction as to which representation should be chosen for each scenario. Conformity of subjects' choices to these predictions provides an indication of subjects' tacit knowledge of the structural properties of the three spatial diagrams. That is, subjects must be able to recognize the structural properties embedded in the story lines, but they need not be able to act upon the properties as objects of reflective thought.

Additional important data come from subjects' verbal justifications for the representations they selected and failed to select. In contrast to the choice data, these data do provide an indication of subjects' explicit knowledge. In the next two subsections, we consider the choice data and the justifications data in turn. Then we present an analysis of the structure of students' knowledge of the three spatial diagrams based on a factor analysis of the justifications data. Preliminary analyses revealed no striking differences between the math educators and computer science majors, so their data were combined in all analyses for increased power.

Representation choice data. Overall, the 23 subjects overwhelmingly chose the type of representation predicted by our structural analysis, with an average proportion of .88 of the choices conforming to our predictions across the 18 scenarios. This level of performance clearly is reliably above chance (.50), $t(22) = 27.90$, $p < .001$, $SE = .0135$. A generally high level of conformity to the predictions of the structural analysis was observed for each representation and structural property, as shown in Table 8. Tests against chance responding (see Table 8) confirm that the support for our predictions generalizes across representations and properties.

However, the data in Table 8 also suggest that the support for our predictions is somewhat weaker for the network representation and for the item/link constraints property. The lower performance in these two categories is due to a single scenario that focused on the network value for the item/link constraints property and had the matrix as the contrasting representation choice. Contrary to our prediction, subjects generally chose the matrix for this scenario, $M = .78$, $p < .02$ by a binomial test. They described constructing a matrix in which the nurses' names labeled both the rows and the columns rather than the network our structural analysis predicts, in which each nurse is represented by a single node. If we exclude this scenario, conformity to the predictions increases to .80 for the network representation and .91 for the remaining two item/link constraints scenarios. We consider the implications of this discrepant result under Discussion.

Justifications data: Overview. We now turn to an examination of subjects' verbal justifications for their representation choices. Each scenario provides an opportunity for subjects to demonstrate their knowledge about two of the three spatial diagrams. Our primary interest is in what subjects know about these representations in general rather than in exactly why they chose a particular representation for a particular scenario. Therefore, our analyses of the justifications data were based on each subject's protocol as a whole. To some extent, what subjects said about a particular scenario depended in part on which other scenarios they had already completed (which varied across subjects). For example, for the later scenarios in the booklet, subjects sometimes prefaced their statements with ''like I said before'' or ''like in the earlier scenario about. . . .''

Summed across the 18 scenarios, subjects averaged 87 codes ($M = 86.87$, $SD = 22.15$), of which 46% reflected statements about the diagram structures. These structural codes provide two sources of evidence for evaluating our structural analysis of the three spatial diagrams.

One source of evidence is subjects' statements that they selected particular representations because the scenarios had the property values specified by the structural analysis. As discussed in Appendix B, for three properties—global structure, number of sets, and link type—we sometimes gave negated codes. For example, a negated number of sets code was given when subjects said they did not choose the network because there were two sets of objects. Such statements are consistent with our structural analysis, but they do not provide strong support for it. Continuing with the number of sets example, subjects might know that networks are not most appropriate when there are two sets of objects but not know that they are appropriate when there is only one set of objects. To provide a conservative test of our structural analysis, our primary analyses include only positive statements that a representation has the property value we ascribe to it.

A second source of evidence comes from considering statements that violate the structural analysis. For example, our analysis specifies that the matrix representation is most appropriate when there are two distinct sets of objects, whereas the network representation is most appropriate when there is only one set of objects. Statements in which a subject claims to have chosen the matrix because the scenario describes two sets of objects or the network because the scenario describes one set of objects support our structural analysis. In contrast, statements in which a subject claims to have chosen the network because the scenario describes two sets of objects or the matrix because the scenario describes one set of objects contradict our structural analysis.

Thus, we would like to see many supportive justifications and few contradictory justifications. This is exactly the pattern we observed. In fact, there were only five contradictory justifications. Given that 23 subjects provided justifications for 18 scenarios each, for a total of 414 justifications that received 930 structural codes, this extreme paucity of contradictory statements (0.54%) may be taken as strong evidence in support of our structural analysis.

Justifications data: A property-centered view. Using the justifications data, it is possible to evaluate the validity of the structural analysis both at the level of the properties (i.e., number of sets, link type, etc.) and at the level of the representation-specific values of the properties (e.g., matrices are most appropriate when there are two sets of objects and networks are most appropriate when there is one set of objects). The property-centered view of the data addresses the issue of whether the properties we proposed are in fact important for distinguishing among the three spatial diagrams. We may ask whether there is support for each of the properties shown in Tables 1–3, except for item distinguishability (because this property was rarely mentioned and could not be coded reliably). If other properties also are important for distinguishing among the three spatial diagrams, this analysis will not uncover them. It is important to note, however, that in addition to the structural codes, which were built into the coding scheme *a priori*, we created codes for any other types of statements that were made with nontrivial frequency across subjects. Thus additional structural properties of the representations could have been identified in this way. It turns out, however, that none of the new codes referred to structural properties of the representations, so if there are additional structural properties, it seems that our (relatively knowledgeable) subjects did not know about them either. We consider data bearing on the properties first, followed by data bearing on the property values. The property-value-centered view of the data addresses the more specific issue of whether the representations have the particular values on the properties that we propose.

For the property-centered analysis, each subject received a score of 0 or 1 for each property, depending on whether he/she gave a justification based on that property for any of the three representations (for any of the 18 scenarios). The average proportion of the 10 structural properties mentioned by our subjects in their justifications was .80 (SD = .09). In fact, all 23 subjects

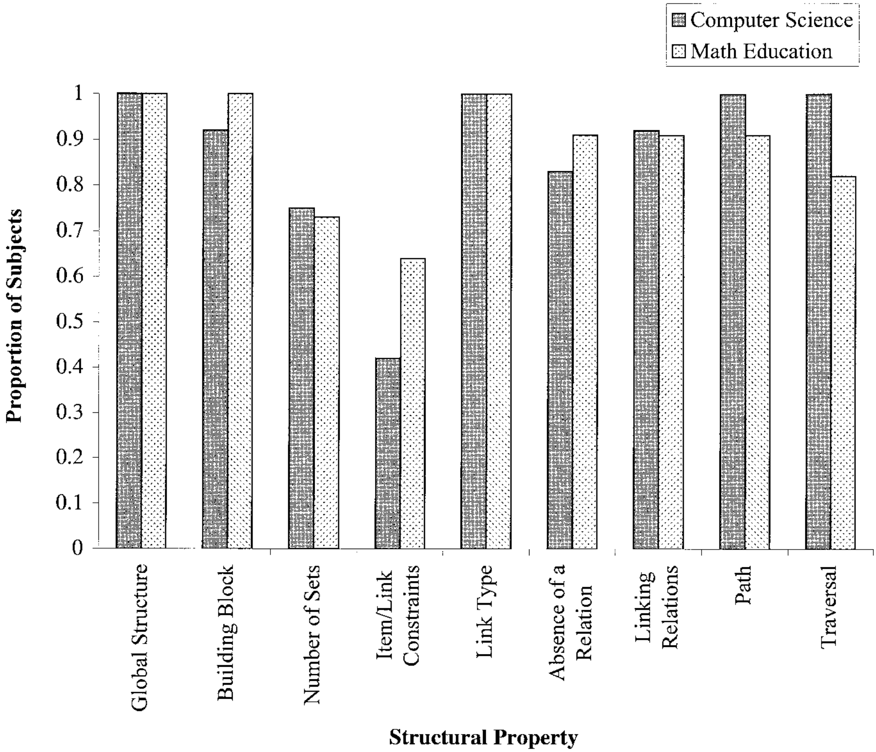

FIG. 2. The proportion of computer science and math education subjects who mentioned each structural property in their justifications for the 18 scenarios.

mentioned a majority of the structural properties in their justifications (range of .60–.90). For the nine properties we were able to code, Fig. 2 shows that the proportion of subjects referring to each property is quite high and similarly so for the computer science and math education subjects considered separately. Considering all 23 subjects together, all nine properties were mentioned by more than half of the subjects, and eight were mentioned by at least 70% of the subjects. This high level of performance provides strong evidence for the properties we hypothesized to be important for distinguishing when to use the three spatial diagrams.

Justifications data: A property-value-centered view. To provide a property-value-centered view of subjects' justifications, Fig. 3 breaks down the property-level data shown in Fig. 2 separately for each representation. In examining these data, it is helpful to set a cutoff for the proportion of subjects who need to mention a property value for a particular representation to count as support for that hypothesized property value. It seems reasonable to select a fairly low cutoff. First, we are essentially looking for an existence proof

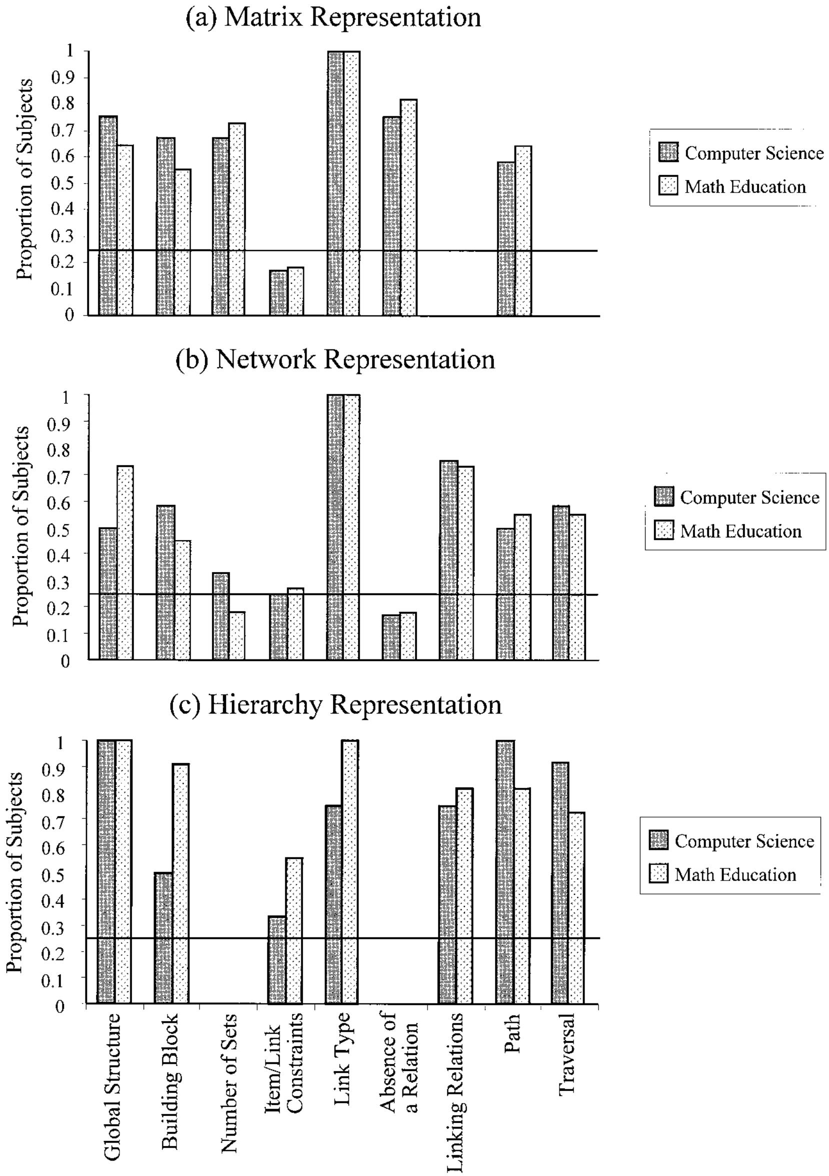

FIG. 3. The proportion of computer science and math education subjects who mentioned each property value for (a) the matrix representation, (b) the network representation, and (c) the hierarchy representation.

TABLE 9
Property Values for Which Positive Evidence Was Obtained from the Verbal Justifications for the Representation Choices for the 18 Scenarios

Property	Matrix	Network	Hierarchy
Global Structure	$+_f$	$+_f$	$+_f$
Building Block	+	+	+
Number of Sets	$+_f$	$+_f$	0
Item/Link Constraints	0_f	$+_f$	$+_f$
Item Distinguishability	0	0	0
Link Type	$+_f$	$+_f$	+
Absence of a Relation	+	0	0
Linking Relations	0_f	$+_f$	$+_f$
Path	+	+	+
Traversal	0	$+_f$	$+_f$

Note. A ''+'' indicates that positive evidence was obtained for a property value (i.e., ≥25% of the subjects spontaneously mentioned that property value). A ''0'' indicates that no evidence was obtained for a property value (i.e., < 25% of the subjects spontaneously mentioned the property value). The subscript ''f'' indicates that the property value was focused in one or more scenarios.

at this point rather than trying to say something about how well known the property values are. Second, we know that a low proportion of subjects mentioning a particular property value does not imply that the remaining subjects made statements that contradicted the structural analysis. As mentioned above, there were essentially no such statements. Third, verbal protocols generally underestimate subjects' knowledge because subjects need not say everything they know (Ericsson & Simon, 1993). Fourth, low cutoffs have been used in other domains in which conclusions have been drawn from subjects' verbal statements (e.g., the category attribute norms collected by Rosch et al., 1976). With these considerations in mind, we selected .25 as our cutoff; that is, at least 6 of the 23 subjects had to spontaneously mention the property value.

As shown in Table 9, 21 of the 30 property values (70%) meet this criterion and therefore receive support from our data. Obviously, there is no support for the item distinguishability property values because we were not able to code that property. For the properties we were able to code, there is supporting evidence for 21 of the 27 applicability conditions (78%).

It is instructive to consider at this point whether adding in the negated property value statements would change the results appreciably. The answer is ''no.'' Negated codes were given for only three properties: global structure, number of sets, and link type. For link type, all subjects who received a negated code for a particular representation also received a positive code, so those results do not change at all. This was also the case for the matrix value of number of sets and the hierarchy value of global structure. The proportions for the matrix and network values of global structure each increase by .17 if negated codes are included, to .87 and .78, respectively (for both subject groups combined). This increase is inconsequential, however, because the proportions for both of these property values based on positive evidence are well above our cutoff.

Where there is some change is for the network and hierarchy values for the number of sets property. Subjects sometimes said that they would not choose the network or hierarchy over the matrix because those representations cannot easily accommodate two sets of objects. If these negated property values are included, the proportion for network increases from .26 to .52 and the proportion for hierarchy increases from 0 to .39 (for both subject groups combined). Both of these increases are sizable, but only the hierarchy increase results in a categorical change. In any case, it is difficult to know whether the negated statements reflect some knowledge about the numbers of sets that can be accommodated most easily in a hierarchy or network or whether they reflect a strong understanding that a matrix is the best representation to use to represent two sets of objects.

Table 9 also indicates, by the subscript ''f,'' which property values were focused in the scenarios. It is clear that the property values for which we did not find supporting evidence primarily were ones that were not focused in any of the scenarios. If we exclude the item distinguishability property, four of the six zeroes were for unfocused properties. Nevertheless, it is not the case that subjects simply parroted back the information we gave them. We found supporting evidence for eight property values that were not focused, and we failed to find supporting evidence for two property values that were focused.

More generally, it is important to note that evidence for particular property values did not only come from the scenarios in which those property values were focused. For the six scenarios that focused on a (single) matrix property value, an average of 4.67 different matrix property values were mentioned (across the 23 subjects) for each scenario. Combining across both the matrix (i.e., the focused) representation and the contrasting representation (either network or hierarchy, depending on the scenario), an average of 5.17 different properties were mentioned for each of these scenarios. For the six scenarios that focused on a network property value, the corresponding means are 6.17 and 6.83. For the six scenarios that focused on a hierarchy property value, the means are 5.17 and 6.33. Collapsed across representations, each scenario received justifications that referred to an average of 5.34 different properties for the focused representation and 6.11 different properties for the focused and contrasting representations combined. In all cases, the maximum number of properties that could have been mentioned is nine. These findings support our initial claim that the 18 scenarios as a whole provided a useful context for eliciting subjects' knowledge about the applicability conditions for the three spatial diagrams.

The Organization of Property Knowledge in Memory

Overview. The data reported in the previous section provide good evidence for our analysis of the applicability conditions for matrices, networks, and hierarchies. However, they do not tell us whether the structural properties are stored in memory as an unorganized list or whether students see the properties as interrelated in some coherent way. This latter possibility is hinted at by the finding that our subjects often mentioned two or more properties in justifying their representation choice for a given scenario. Moreover, our structural analysis suggests that certain properties are more closely related than others. For example, the three ‘‘movement’’ properties are predicted to be highly interrelated: linking relations is the local property that affords the more global differences in traversal, and traversal is a more detailed specification of the types of paths that are possible. Similarly, link type and absence of a relation should be closely related because one can conceptualize representing particular types of relations between pairs of items and representing the absence of any relation between pairs of items as flip sides of the same coin. In addition, global structure and building block should be closely related.

Because it is reasonable to expect that some structural properties may tend to be used together in justifying a representation choice (e.g., linking relations and traversal), whereas other properties may have an antagonistic relationship such that if one property is used in the justification the other property will not be used (e.g., traversal and number of sets), factor analysis seemed to be the most appropriate statistical tool for our purposes. The factor analysis was conducted on nine dependent measures, with each measure being references to one of the structural properties (all except item distinguishability). The replications factor was the different scenarios. Because we know the focused property, the two representations presented, and the representation predicted by the structural analysis to be most appropriate for each scenario, we can use this information to help interpret the results of the factor analysis.

We conducted a single factor analysis that included the data from the math education and computer science students separately. If the results are consistent across the two groups, we can be more confident that they reflect true interrelations among the structural properties. Moreover, this decision gave us 36 cases for the analysis (2 per scenario) rather than only 18, which was desirable because we had nine dependent measures. Thus, the data we analyzed were the average frequencies with which the nine properties were used by each subject group to justify the representation choice for each scenario. For example, the global structure score for Scenario 1 for the math education students indicated the average number of references to that property by those subjects, regardless of whether the coded statements referred to the matrix or to the hierarchy (the two representations presented with that scenario). The average correlation (computed across the 18 scenarios) between the scores of the two subject groups for the nine property variables is .76.

We obtained a four-factor solution that accounted for 72% of the variability in the data. To facilitate discussion and interpretation of the results, only property variables with factor loadings of at least .30 in absolute value are considered. To interpret the factors, we examined each observation (i.e., a particular scenario for a particular subject group) to determine what features were shared by those with high versus low scores on each factor. In almost all cases, the factor scores for a particular scenario for a particular factor were quite similar for the two subject groups.

Interpreting Factor 1. Figure 4a shows the results for Factor 1, which accounted for 29% of the variance. Linking relations and traversal loaded positively on this factor, whereas number of sets, link type, and absence of a relation loaded negatively. Looking at the factor scores for the scenarios, three distinct groups emerged. Large positive factor scores were obtained by the four network versus hierarchy scenarios that focused on either linking relations or traversal (four scenarios times two subject groups yields eight observations). The remaining scenarios that contrasted the network and hierarchy representations had factor scores close to zero. All scenarios in which one of these two representations was contrasted with the matrix representation had moderately large negative factor scores. The factor scores were somewhat more extreme when matrix was the correct representation than when it was the contrasting representation (*M*s of −.61 and −.39, respectively). These differences between the scenarios also are evident in the property (i.e., component) scores. The first group of scenarios had many linking relations and traversal justifications but few number of sets, link type, or absence of a relation justifications. The scenarios in which matrix was one of the representation choices showed the opposite pattern of justifications. The remaining network versus hierarchy scenarios—those that were not written to focus on either linking relations or traversal—showed an intermediate pattern.

These results suggest that a key decision in selecting an appropriate representation is whether to use the nonmovement representation (the matrix) or one of the two movement representations (the hierarchy or network). If the latter choice seems most prudent, then the number of allowable connections into and out of a node (linking relations) and the implications of this for traversing the representation are used to determine whether the network or the hierarchy best captures the structure of the scenario. Thus linking relations and traversal justifications are given when these two representations are contrasted, regardless of which one is most appropriate. As indicated in Table 3, these two properties are critical for distinguishing the hierarchy and network representations. Note, however, that subjects primarily accessed this knowledge for the four network versus hierarchy scenarios that focused on one of these two properties.

At the other end of Factor 1, number of sets, link type, and absence of a relation combine to distinguish the hierarchy and network representations from the matrix representation, regardless of which representation best cap-

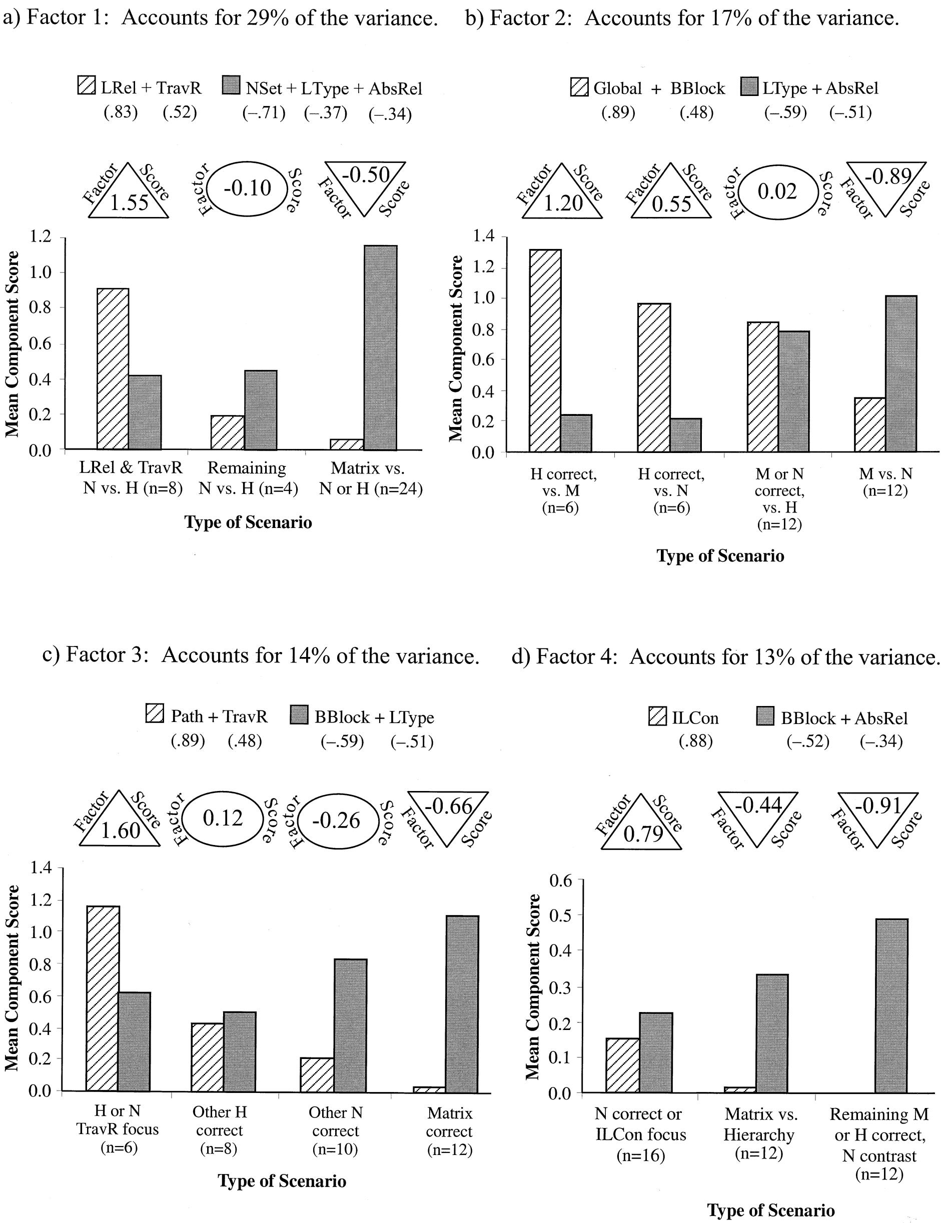

FIG. 4. (a) Factor 1 scores and associated property component scores for distinct sets of scenarios. (b) Factor 2 scores and associated property component scores for distinct sets of scenarios. (c) Factor 3 scores and associated property component scores for distinct sets of scenarios. (d) Factor 4 scores and associated property component scores for distinct sets of scenarios. In all four panels, factor loadings for the property variables are indicated in parentheses below the property names at the top of the figure.

tures the structure of the scenario. These three properties describe important ways in which a hierarchy or network differs from a matrix (see Tables 1 and 2): Hierarchies and networks, but not matrices, can represent directional relations. Matrices, but not the other two diagrams, can explicitly represent the absence of a relation. And matrices and networks differ in the number sets of objects they are optimized to represent.

Interpreting Factor 2. Figure 4b shows the results for Factor 2, which accounted for 17% of the variance. Global structure and building block loaded positively on this factor, whereas link type and absence of a relation loaded negatively. Looking at the factor scores for the scenarios, four distinct groups emerged. Large positive factor scores were obtained by the three scenarios in which the hierarchy and matrix representations were contrasted, and the hierarchy best fit the structure of the scenario. Smaller positive scores were obtained by the three remaining scenarios that best fit a hierarchy representation, in which network was the contrasting choice. The scenarios for which hierarchy was one of the two choices, but that representation did not best fit the given structure, had factor scores of essentially zero. Finally, scenarios that contrasted the matrix and network representations, regardless of which one best captured the structure of the scenario, had large negative factor scores. These differences between the scenarios also are evident in the property (i.e., component) scores: The sum of global structure and building block justifications decreases across the four categories of scenarios, whereas the sum of link type and absence of a relation justifications increases.

The primary function of this factor seems to be to distinguish the hierarchy from each of the other two representations and secondarily to distinguish the matrix and network from each other. The results suggest that the global structure and building block properties together are important for discriminating the hierarchy from the matrix and the network. These justifications were most common when the hierarchy best fit the structure of the scenario and that representation was contrasted with the matrix, but they occurred with considerable frequency whenever hierarchy was one of the representation choices. This makes sense. The hierarchy and matrix representations both have clearly identifiable structures, so determining that the current situation has a levels or contingency structure is a strong cue both for the hierarchy and against the matrix. In contrast, because the network is the more general representation, such information does not provide strong evidence against it. In fact, the network can be used to mimic either of the other two representations, and our subjects clearly knew that with respect to the hierarchy. In justifying their choices for the hierarchy correct, network contrast scenarios, 61% of the subjects indicated that a network can be used to mimic a hierarchy (this was one of our nonstructural codes).

At the other end of Factor 2, link type and absence of a relation together are important for distinguishing the matrix and network representations from each other, regardless of which one is most appropriate. This makes sense because each of these properties specifies that one of these two representations can do something important that the other cannot: The network but not the matrix can represent directional relations, and the matrix but not the network can represent the absence of a relation.

Interpreting Factor 3. Figure 4c shows the results for Factor 3, which accounted for 14% of the variance. Path and traversal, which loaded positively on this factor, are global properties of the representations (i.e., properties that are only evident when one considers the representation as a whole); whereas building block and link type, which loaded negatively, are local properties. Looking at the factor scores, four distinct groups of scenarios emerged. Large positive factor scores were obtained by the three scenarios in which the hierarchy or network value of traversal was the focused property. Two of these scenarios contrasted the hierarchy and network representations. For the third scenario, hierarchy was the correct representation and matrix was the contrasting representation. The remaining scenarios for which hierarchy or network provided the best structural model had small positive or negative factor scores, respectively. Finally, all scenarios for which the matrix provided the best structural model had moderately large negative factor scores. These differences between the scenarios also are evident in the property (i.e., component) scores: The sum of path and traversal justifications decreases across the four categories of scenarios, whereas the sum of building block and link type justifications increases.

These groupings of scenarios make sense given this factor's focus on the global/local dimension. Consider the distinction between the scenarios focusing on the hierarchy or network values of the traversal property and the other scenarios for which one of these two representations provides the best structural model. Hierarchy and network are the representations of choice for illustrating paths, and traversal is a key property distinguishing when to use a network versus a hierarchy. Thus, in many situations, these two representations are chosen because of their ability to represent certain kinds of paths. In other situations, however, these representations are chosen because they are able to represent local relations of certain types or because they are composed of a particular type of structural unit. For example, a power hierarchy can be used to illustrate the chain of command and thus to determine the sequence of people who must approve an idea generated by someone at the lowest level. This same hierarchy, however, could also be used to determine who is under the command of a particular person. The justifications for choosing a hierarchy are likely to differ depending on whether the goal requires attention to global or local information. Similarly, a bus route map (network) can be consulted to plan an extended trip with many stops along the way or to determine whether it is possible to travel nonstop from City X to City Y. Again, the justifications for choosing a network in these two cases are likely to differ.

In our study, subjects gave global justifications for the scenarios that highlighted the need to represent such information, i.e., those hierarchy or network correct scenarios that focused on the traversal property. For all three of these scenarios, both traversal and path justifications were given by at least several subjects from each subject population. The other hierarchy and network correct scenarios received some such justifications, but they also received local justifications appropriate to the purpose for which spatial diagrams are useful for those scenarios.

Finally, consider the scenarios for which the matrix provides the best structural model. These six scenarios essentially received only local justifications, as is appropriate. Matrices are not useful for representing paths of any sort. However, they are very good for indicating which pairs of items are associated or for cataloging characteristics associated with pairs of items.

Interpreting Factor 4. Figure 4d shows the results for Factor 4, which accounted for 13% of the variance. Item/link constraints loaded positively on this factor, whereas building block and absence of a relation loaded negatively. Looking at the factor scores, three distinct groups of scenarios emerged. Positive factor scores were obtained by the scenarios for which the network representation was most appropriate and by the scenarios that focused on the item/link constraints property. The scenarios that contrasted the matrix and hierarchy representations had moderately negative factor scores. The remaining scenarios for which the matrix or hierarchy representation provided the best structural model and the network was the contrasting representation had sizable negative factor scores. These differences between the scenarios also are evident in the property scores: Only the first group of scenarios received justifications based on the item/link constraints property. Note, however, that such justifications were not prevalent even for these scenarios. As we discussed above, this is not a common justification. The sum of the building block and absence of a relation justifications increased across the three categories.

This factor seems to focus on the question of whether there are constraints on which items can be connected. If the answer is no, then the best representation to use is the network. A key defining feature of the network is that the absence of any formal predefined structure (global structure value) means that any item can potentially be connected to any other item (item/link constraints value). Thus, the item/link constraints property plays an important (although admittedly small) role in justifying one's representation choice when the network representation provides the best structural model for the scenario. Subjects also considered this property when network was the contrasting representation and item/link constraints was the property focused.

In contrast, if there are constraints on which items can be connected, then either the matrix or the hierarchy representation would be the best one to use. When a choice of one of these two representations must be justified over the network representation, then the presence versus absence of constraints is

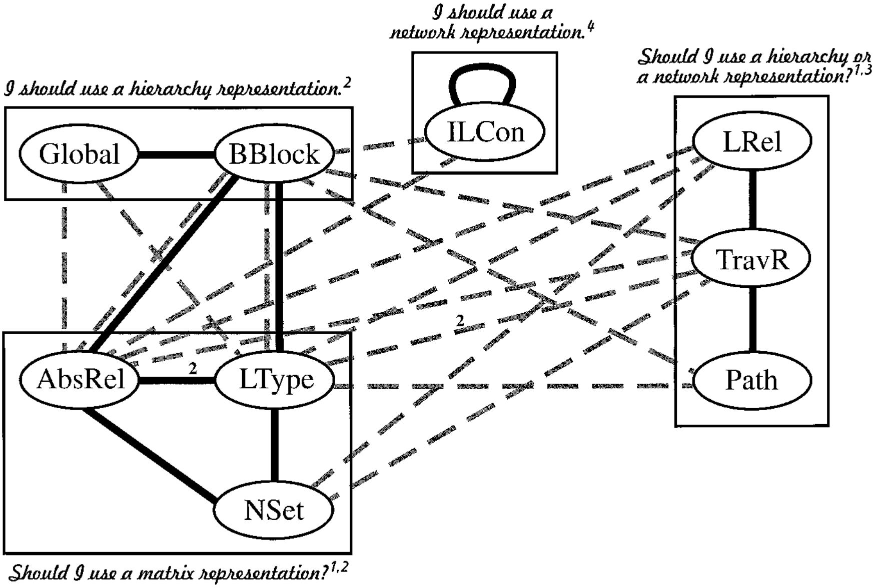

FIG. 5. Summary of the factor analysis results. Solid black lines indicate property variables that loaded positively on the same factor. Dashed gray lines indicate property variables that loaded negatively on the same factor. Superscripts on the annotations indicate which factors provide evidence for the indicated property groupings. From ''Spatial diagrams: Key instruments in the toolbox for thought,'' by L. R. Novick, 2001, *The psychology of learning and motivation* (Vol. 40, pp. 279–325), D. L. Medin (Ed.), San Diego: Academic Press. Adapted with permission.

highlighted, and large negative factor scores are observed. For the hierarchy representation, subjects tended to talk about these constraints in terms of which particular items are connected below a certain other item (building block property; consistent with the positive end of Factor 2), whereas for the matrix representation they tended to talk about the constraints in terms of certain items not being paired (absence of a relation; consistent with the negative end of Factor 2). Scenarios in which the hierarchy and matrix were contrasted yielded intermediate factor scores, presumably because the presence or absence of constraints was not a very useful factor to consider in making one's representation choice (because constraints are present for both).

Summary of the factor analysis results. Figure 5 summarizes the relations among the properties that are indicated by the results of the factor analysis. In this figure, properties that loaded on the same factor with the same sign are connected by a solid black line. Properties that loaded on the same factor but with the opposite sign are connected by a dashed gray line. This annotated network representation of the factor analysis results helps to clarify the underlying organization of the properties.

Consider, first, the right side of the figure. As we predicted, linking relations, traversal, and path form a coherent cluster. Traversal is positively related to both path and linking relations. As we discussed above, the traversal property is a more detailed specification of the types of paths that are possible, and differences in the local linking relations property are responsible for the more global differences in traversal. These three properties are perceived as being quite distinct from the remaining properties, as the remaining 10 links coming out from those nodes represent negative relations (note that there are two negative links between traversal and link type). Moreover, those negative links go to four of the six remaining properties (all except global structure and item/link constraints). As we noted in our discussion of Factors 1 and 3, this cluster of properties is most important for distinguishing the hierarchy and network representations.

Absence of a relation, link type, and number of sets also form a fairly coherent cluster. Each of these properties was positively linked to each of the other properties at least once in the factor analysis, and there were no negative relations among these properties. Link type and absence of a relation were especially strongly related, as we expected. This cluster of properties is most important for determining whether a matrix would be the most appropriate representation, as we discussed in conjunction with Factors 1 and 2. This cluster of properties is not as distinct as the linking-relations/traversal/path cluster, however, because both absence of a relation and link type have positive as well as negative relations to building block. We expected the positive relation between the latter pair of properties because link type specifies the type of relation in the building block. These two properties were used together in justifications when a matrix provided the best structural model for the scenario (Factor 3). The positive link between absence of a relation and building block does not have as coherent an explanation, as discussed above in connection with Factor 4.

Global structure and building block form a third, smaller cluster, as expected. As discussed in conjunction with Factor 2, these properties are important for determining that a hierarchy representation would be more appropriate than either of the other two representations.

Finally, item/link constraints forms a cluster all to itself, as indicated by Factor 4. This property was used almost solely to justify the choice of a network representation. We expected that item/link constraints might be grouped with global structure because the global structures of the representations are determined in part by the constraints on which items can be linked. We discussed the interdependence between these two properties for the network representation in our interpretation of Factor 4. Although there is no direct support for such an interdependence from our computer science and math education subjects, there also is no evidence against this hypothesis. Interestingly, the results of a factor analysis of the justifications given by a group of more typical undergraduates indicated that global structure and item/link constraints were grouped together for those subjects (Novick, 2001).

It is important to note that Fig. 5 and the factor analysis on which it is based represent property relations collapsed across representations. It is possible that there are some exceptions or qualifications to these relations for individual representations. We consider this issue in more detail in the Discussion, as our present data do not allow us to address it empirically.

Summary of the Empirical Findings

The results we have presented here allow us to draw several conclusions concerning the validity of our structural analysis of matrices, networks, and hierarchies. Most importantly, we found strong support for our structural analysis with respect to both subjects' choices of the best representation for each scenario and their verbal justifications for their choices. Subjects overwhelmingly selected the representations predicted by the structural analysis (88% conformity to the predictions). This conclusion holds for the individual representations as well as for the individual focused properties. This shows that subjects' implicit knowledge about which representations are useful for different types of situations is consistent with our analysis of the representations' structures.

In examining subjects' verbal justifications for their selections, we found even more direct support for our structural analysis. We were able to reliably code 9 of the 10 properties we proposed in our analysis of the representations' structures. Subjects either did not know the item distinguishability property or did not consider that property to be relevant for performing our task. For the 9 codable properties, subjects almost never made statements contradicting the structural analysis (e.g., saying that the network is most useful when there are two sets of objects rather than one). Moreover, all subjects mentioned a majority of the 10 structural properties somewhere in their protocols ($M = 80\%$). It is not surprising, therefore, that all 9 codable properties were mentioned by more than half of the subjects, and 8 of the 9 were mentioned by at least 70% of the subjects. Overall, 21 of the 27 codable property values (9 properties $\times$ 3 representations; 78%) were supported by evidence from the verbal justifications provided by our subjects. (Two of the remaining six property values were mentioned by only a handful of subjects, and the remaining four were not mentioned at all.) These results indicate that the students' explicit knowledge about the applicability conditions for the three spatial diagrams supports our analysis of the structures of those representations.

As we discussed in the introduction, the structural properties we proposed are not orthogonal. Rather, various of the properties are conceptually related. The results of our factor analysis, summarized in Fig. 5, indicate that the properties are organized in memory in a coherent manner that is consistent across subjects. Moreover, many of the observed relations between the properties were expected based on our structural analysis of the three representations.

As we predicted, linking relations, traversal, and path—the three movement-related properties—formed a coherent cluster that was used to determine whether a network or a hierarchy representation provided the best structural model for a scenario. This cluster is clearly distinct from the other properties, suggesting that the distinction between static and dynamic aspects of the representations is an important one. The global structure and building block properties grouped together in another cluster, as we predicted. These properties were especially important for determining that it would be good to use a hierarchy representation. The item/link constraints property formed its own separate cluster and was primarily relevant for determining that a network representation would be appropriate. A final cluster contained the link type, absence of a relation, and number of sets properties. These properties were most important for determining whether to use a matrix representation. The association between link type and absence of a relation, which both concern local properties of the representations, was expected. As indicated in Table 1, however, we had thought of number of sets as a more global property of the representations.

Finally, there were positive links between the building block property and both the link type and absence of a relation properties. The association between building block and link type was expected. However, there were also negative links between building block and these other two properties, stemming from building block's association with the global structure property. Building block seems to play a dual role in selecting an appropriate spatial diagram, sometimes supporting local properties (especially when the focus is on the matrix representation) and other times supporting the global structure property (when the focus is on the hierarchy representation).

DISCUSSION

Contribution of the Present Research

Existing structural analyses of visual representations (Lohse et al., 1994; Twyman, 1979) indicate that matrices, networks, and hierarchies have different structures, but they are not detailed enough to guide specification of the applicability conditions for these representations. We propose that matrices, networks, and hierarchies differ with respect to 10 interrelated structural properties. The properties are like slots in a schema, and the (nominal) values for these slots constitute the applicability conditions for the representations. Although previous theoreticians have proposed individual structural features of these representations, there has been little discussion of the relations among the features or the representations. Thus, our structural analysis as a whole is new. Moreover, all of the previous work has been theoretical rather than empirical, so there is no empirical evidence that the property values are actually used as applicability conditions for determining which spatial diagram best fits the structure of the situation under consideration.

Our subjects' representation choices and their verbal justifications for those choices provide good support for 9 of the 10 structural properties we proposed: global structure, building block, number of sets, and item/link constraints (see Table 1); link type and absence of a relation (see Table 2); and linking relations, path, and traversal (see Table 3). In addition, the results of our factor analysis of the justifications data support many of the expected interrelations among the properties. These results are important in two respects. First, they validate our characterization of three diagrammatic tools that have been shown to be important for reasoning, problem solving, and learning. Second, they indicate that the property values have psychological force because subjects used them to decide on and justify their choice of the best representation for each of our scenarios. In fact, nearly half of the codes (46%) we gave to the verbal protocols reflected the proposed structural properties (i.e., applicability conditions) of the representations. Another 23% of the codes reflected subjects' use of the three spatial diagrams as representations by associating specific terms from the scenarios with specific parts of the diagrams.

In the remainder of the Discussion, we consider possible revisions to our structural analysis that might be required as well as some of the limitations of the present findings. We also discuss the relation between structure and function as well as the source of students' knowledge about the three spatial diagrams. Consideration of these issues will lead us to propose some directions for future research.

Possible Revisions to our Structural Analysis of the Three Spatial Diagrams

The status of property values that were not supported. Nine of the 30 property values (10 properties $\times$ 3 representations) were not mentioned by a critical number (at least 25%) of subjects. Three of these property values were for the item distinguishability property (see Table 2), which we were not able to code reliably (in part due to the paucity of references to it). The remaining six property values were (a) matrix item/link constraints, linking relations, and traversal; (b) network absence of a relation; and (c) hierarchy number of sets and absence of a relation.

What should we conclude from the absence of support for these nine property values? In addressing this issue, it is important to keep in mind two aspects of our experimental design and results. With respect to experimental design, we should note that we did not pick a random sample of property values on which to focus in the scenarios. Rather, we chose those that we thought were most important or most discriminating. With respect to our results, we reiterate that although our subjects did not spontaneously mention these nine property values, they also did not make statements that contradicted them. Thus, the status of these property values remains uncertain. This uncertainty could end up being resolved in a number of different ways.

The most unfavorable possibility is that the property itself is unimportant for distinguishing when to use each of the three spatial diagrams. This possibility is relevant only to the item distinguishability property because the remaining properties were mentioned by a critical number of subjects for at least one of the representations. We suspect that the item distinguishability property may be important, at least for discriminating between the network and hierarchy representations, but that this did not show up in our results for any of the reasons outlined in the subsequent paragraphs.

A second possibility, which is more generally relevant, is that the uncertain property values are important in some situations, but they were not needed to justify the representation choices for our scenarios. The fact that seven of the nine uncertain property values were not focused on in any of our scenarios is consistent with this possibility. On the other hand, as we discussed under Results, the relation between focusing on a property value and finding evidence for it was not a simple one. Thus it is difficult to evaluate the plausibility of this explanation. We suspect, however, that it may apply in several cases, particularly network absence of a relation and network and hierarchy item distinguishability and perhaps also matrix item/link constraints.

The third possibility is that not all of the properties provide important applicability conditions for all of the representations. As discussed in the introduction, we forced our structural analysis to have a factorial structure, in which each of the properties specified an applicability condition for each of the representations. It may be the case that this factorial structure is only a useful approximation. Looking at Tables 1–3, some property values seem less diagnostic than others. These property values may include (a) matrix item distinguishability, linking relations, and traversal; and (b) hierarchy number of sets and absence of a relation.

If our intuitions are correct, then the number of properties that provide important applicability conditions varies across representations. In particular, we suspect that only the network representation may have an applicability condition corresponding to each property. It makes sense that the properties important for specifying the matrix and hierarchy representations may be a subset of those important for specifying the network representation because, mathematically, the matrix and hierarchy are special cases of the network (e.g., Chartrand, 1985). These special cases are created by altering properties that specify the general form.

Further consideration of the building block property. Overall, we obtained very good reliability for coding the structural properties from the verbal protocols. However, we did encounter some difficulty, initially, in coding the matrix value of the building block property because that code was sometimes confused with the matrix link type and global structure codes. These two properties are ones with which building block is highly associated, both in our structural analysis and empirically. It is possible, however, that this coding difficulty reflects a more general underlying deficiency in students' representational knowledge. Regardless of the representation being discussed, subjects typically mentioned the building block entirely embedded within the context of the cover story. They very rarely explicitly stated the idea that the building block is repeated a number of times to construct the representation. Thus it might be wise to consider support for the building block property to be qualified at this point. Alternatively, one might conclude that students' ability to use this property is particularly rudimentary.

On the possibility of additional applicability conditions. A great strength of the verbal protocol methodology is its ability to uncover relevant information that the researcher has not considered. Thus, we are encouraged that our protocol analysis yielded no additional structural properties used by subjects to guide their representation choices. However, we did uncover a size constraint on the network value for the number of sets property. This constraint was revealed in subjects' responses to Scenario 9, for which they generally chose the matrix representation rather than the network representation predicted by our structural analysis. The scenario is as follows:

> A new study indicates that nurses who have worked together in the past are more efficient than nurses who have not worked together. Due to changing priorities at the hospital, the head nurse must change the current work assignments. Although she is free to assign any of the nurses to work together that she wants, the head nurse places a high value on efficiency. Therefore, she would like a diagram showing which nurses have worked together in the past.

We proposed that a network is most appropriate when the represented world contains only a single set of objects. A number of subjects (mostly from the computer science population) indicated that this condition holds when the number of objects to be represented is relatively small. When the number of objects is large, however, the network is likely to be unwieldy and therefore a matrix is preferred. These subjects said they chose the matrix representation because there might be a lot of nurses, in which case it would be easier to find a particular nurse's name in a matrix than in a network. They implied that if they knew that there were not too many nurses, they would have chosen the network representation.[6] Subject 13 gave an especially clear statement of the size constraint:

> You could model you could put all of the nurses um in the rows and all of the nurses in the columns. Let's see. And show specific interactions in between. This might

[6] These students were not simply reporting what they had been taught in their computer science courses. For example, the textbook currently being used for the introductory data structures course at Vanderbilt University (Weiss, 1999) indicates that, although, in general, networks can be represented as two-dimensional arrays, the space requirement for the matrix can be prohibitive if there are very many nodes in the network and only a small proportion of them are connected (because most of the matrix cells would contain a zero).

> be a little inefficient because one nurse will never, I mean going down the diagonal you'll never actually have a nurse interacting with herself or working with herself. Um, but the network, if there are a lot of nurses this could get really confusing um with paths just going all over the place. Um. For a small group the network might be a little bit more efficient. But for a large group which is probably what you've got in a hospital, then uh the matrix might be a little bit bigger might be a little more inefficient spacewise. But it would easier for the person to see exact interactions.

Bertin (1981, p. 129) is explicit on this issue, stating that the network ''ceases to be a means of discovery when the elements are numerous. The figure rapidly becomes complex, illegible and untransformable.'' With 20/20 hindsight, we concur.

This discussion highlights the need to consider factors that influence the ease of using the various representations. We will consider this issue in a later section.

Further Consideration of the Structural Properties and Their Interrelations

In addition to identifying the structural properties that are used to determine when one spatial diagram versus another is most appropriate, it is important to consider whether the diagnosticity of the properties varies across representations. For example, we suspect that the number of sets property is more important than the traversal property for specifying a matrix representation, whereas the reverse is the case for the hierarchy representation. To take another example, although subjects mentioned the path property for all three representations, we suspect that it is more important for specifying a network or hierarchy than a matrix. It may also be the case that the links between the properties and the representations are asymmetric. For example, the link from the matrix representation to the matrix value of the linking relations property may not be equivalent in strength to the link from that property value to that representation. Although a matrix is easily applicable to a situation in which both one-to-many and many-to-one relations must be represented, knowing that the represented world contains both one-to-many and many-to-one relations may not strongly cue the use of a matrix. It is not possible to get information about the strengths of the associations (in either direction) between the structural properties and the representations from our data.

A second important issue to consider concerns the nature of the interrelations among the structural properties. Some properties are likely to be closely related for all three representations: for example, (a) global structure and building block and (b) building block and link type. For other pairs of properties, however, the nature of the relation between them may depend on the representation. For example, traversal may be closely related to linking relations and path for the network and hierarchy representations but not for the matrix representation.

Figure 5 shows relations that generalize across representations but not those that are specific to individual representations because our factor analysis was conducted on data that were collapsed across representations. To discover representation-specific relations, we would have to conduct factor analyses of the property variables separately for each representation. Two considerations made it inadvisable to conduct such analyses. First, each representation was used with only 12 of the 18 scenarios, so the number of cases in the analyses would have been quite small. Second, not all of the properties were mentioned for all of the representations (see Fig. 3). It is difficult to determine whether the relation between properties X and Y differs across representations unless both of those properties were used for all three representations.

The Need for Varied Methodologies

We learned a lot about the applicability conditions for the three spatial diagrams from the representation-choice and verbal-protocol methodologies used in the present study. A strength of the choice methodology is that it allowed us to tap into representational knowledge that might be to some extent implicit, whereas the verbal protocols required explicit access to representational knowledge. One strength of the protocol methodology is that it yielded quite strong evidence for the structural properties that subjects did in fact use in their representation justifications. Another strength of this methodology, mentioned above, is the possibility of discovering additional properties that were not proposed initially.

Nevertheless, there are limitations to the kinds of answers these methodologies can provide. We have indicated several issues that cannot be resolved with our data: (a) the status of the property values that were neither supported nor refuted, (b) the diagnosticity of each property for specifying the use of each representation, (c) the symmetry or asymmetry of the links between the properties and the representations, and (d) the extent to which relations among the properties depend on the representation. A major stumbling block in addressing these issues is that verbal protocols are informative with respect to what they contain but not with respect to what they do not contain. It is difficult to make good inferences from what subjects fail to say. Thus, to address all of these issues, it will be valuable to switch to a more sensitive, close-ended response method (e.g., asking subjects to generate some type of rating) for subsequent studies.

The Source of Students' Knowledge about the Three Spatial Diagrams

Induction over ordinary life experiences. Our subjects clearly were quite knowledgeable about matrices, networks, and hierarchies. It is important to consider the source of their knowledge about these representations. We have argued previously, based on other research, that students possess rudimentary abstract schemas for the three spatial diagrams (Hurley & Novick, 2001) and that the content of these schemas has been induced over ordinary life experiences (Novick, 2001; Novick et al., 1999). That is, the examples of these representations in use that people encounter in everyday life provide sufficient information for them to induce some of the structural properties of the representations. In U.S. culture, examples of these representations in use are very common. For example, they appear in newspapers and magazines, and they are seen in science textbooks beginning in elementary school. Although the representations per se are not the focus of study in school (see the next subsection), with repeated use of these representations in the service of other goals (e.g., learning about the food web for a particular ecosystem), students have the opportunity to learn something about the representations themselves. We propose that students avail themselves of this opportunity, at least to some degree.

It is interesting to consider whether induction of abstract knowledge about matrices, networks, and hierarchies is confined to modern Western culture. Although we are not aware of any cross-cultural research on this issue, we believe that the capacity for abstract diagrammatic representation is a general human trait. There is evidence for the use of spatial diagrams, particularly matrices and networks, for various purposes among many disparate cultures going back several millennia.

Magic squares (square matrices in which numbers in the cells sum to a constant across rows, columns, and the main diagonals) originated in China, where myth indicates that they date back to about 2200 BCE (D. Smith, 1973; Zaslavsky, 1973). Their first known appearance outside of China is in Arabia in the 9th century. Magic squares appeared in India around the 11th to 12th centuries and were introduced to Europe around the 14th century. Multiplication tables printed in matrix form have been known since the 15th century (D. Smith & Ginsburg, 1956). Matrices have also been used as game boards in various cultures (Zaslavsky, 1998). In several countries in northwestern Africa, people play a game on a 5×6 board in which counters are moved according to specific rules in an attempt to get three counters in adjacent cells in a row or column. This same board is used in West Africa to play a game similar to checkers.

Networks also have a long and cross-cultural history. A painting of what looks like a route map, with abstractly represented buildings connected by paths, has been found on a rock in Ponte San Rocco, Italy. This petroglyph of a network in use dates from the middle or late Bronze Age, ca. 1900–1200 BCE (C. Smith, 1987).

Abstract networks, like the one shown in Fig. 1, are common in many cultures (although not in the United States). For example (Zaslavsky, 1998), cultures from around the world (e.g., several African countries, The Philippines, England, Columbia, Sri Lanka, and Pueblo Indians in North America) play games in which players take turns moving their counters along the links in a specific network (the network varies across cultures) until one player gets three counters in a row (i.e., on three nodes directly connected along a path). Game boards for several of these games were found carved on the roof of an ancient Egyptian temple, ca. 1300 BCE. In Mozambique, India, and Bangladesh, people play a checkerslike game on a butterfly-shaped network.

Early in the 12th century, a Belgian visitor to the Congo Basin in Zaire introduced European mathematicians to the networks used by the Shongo people who live there (Braxton, Gonsalves, Lipner, & Barber, 1995; Zaslavsky, 1973). Shongo children play a game in the sand in which they trace complex networks (e.g., one network has 142 nodes and 251 links)—every link in the network is traversed exactly once without lifting one's finger from the sand. Networks also play an important role in Shongo culture as part of their storytelling and oral traditions. As these people tell stories, they draw networks to represent the events and changes that take place.

Ruling out alternative hypotheses. In this section, we consider and rule out three alternative hypotheses concerning the source of our subjects' knowledge about matrices, networks, and hierarchies. One hypothesis is that our subjects induced the applicability conditions for these representations from examining the diagrams we presented as the choices for the selection task (see Fig. 1). But if subjects are able to do this induction so easily, surely they would have done so based on their previous experiences with these diagrams (also see Novick et al., 1999). As noted by Novick et al., even unselected samples of college students have many such experiences.

A second hypothesis is that perhaps students were taught how to reason with these diagrams in high school math classes or enrichment programs. The first author has asked undergraduates this question in the context of teaching seminars on thinking with diagrams. Although the students indicated that they had encountered matrices, networks, and hierarchies in various contexts both inside and outside of school, they had not been directly taught about the representations. That is, the representations themselves were never the object of study. The only exception is that a small handful of students said they learned how to use the matrix representation to solve certain kinds of deductive reasoning problems. Even these students, however, were not taught structural properties of that representation.

A third hypothesis is that our subjects simply told us what they had been directly taught in the college courses in their majors. We believe that this is not the case for either of our subject groups. Looking at the list of suggested mathematics courses for teacher licensure candidates, it is very unlikely that our math education students had taken courses in which the three spatial diagrams were the focus of study.

Although the computer science students had taken such courses (that was why we selected them), it is definitely not the case that the textbooks present our structural analysis of any of the three representations. We examined four textbooks (Cormen et al., 1990; Elmasri & Navathe, 1994; Epp, 1995; Weiss, 1999) that are currently used in courses that our computer science subjects would have taken prior to participating in our study. Although we do not know that our subjects used these particular textbooks, they would have used comparable books. From these four books combined, we found references to 14 of our 30 proposed property values (12 of the 21 for which we found support in our study). These references were scattered across chapters and across textbooks, which were encountered in courses taken over at least 2 years. No textbook made reference to more than 7 property values overall or to more than 4 property values (of 10 possible) for any one representation. Moreover, in many cases the structural features were presented not as abstract properties of the representations but as characteristics of certain situations that can be modeled in a certain way in a computer program.

It is also relevant to note that relations among the features were not discussed in any of the textbooks, yet our subjects clearly saw the structural properties as being related. Moreover, the perceived relations were consistent across subjects, both within and across the computer science and math education groups.

Representation Structure, Function, and Use

As we come to the end of our discussion, we revisit the relation between structure and function for the three spatial diagrams. Our consideration of the functions of these representations brings into focus issues concerning the use of these representations for reasoning and problem solving, a topic that so far has remained in the background. We provide a brief discussion of this topic here.

We argued earlier that the primary function of matrices, networks, and hierarchies is to convey problem structure, so we consider our analysis of the properties of these representations to be structural in nature. For example, if one's goal is to represent the absence of a relation between two items, the structure of the matrix is better suited than that of the hierarchy or network to accomplish this function. In contrast, if one's goal is to represent paths connecting multiple items, the hierarchy or network would provide a better structure than would the matrix to accomplish this function. Our perspective on this issue is consistent with that of others who have argued more generally that the function of visual representations is to convey structure (Lohse et al., 1994; Sanfey & Hastie, 1998; Winn, 1989).

However, it is likely also the case that because representations have different structures, they make certain kinds of inferences easier or more difficult to compute. Put another way, different structures may afford different types of inferences. This may be especially true for some of our movement properties (e.g., path and traversal). Our claim is only that the structure comes first, which is why we set out to investigate representation structure in our research.

Nevertheless, we do not wish to give the impression that the inferences afforded by the different representations are unimportant. In fact, there is evidence that they are quite important for using these representations in the service of reasoning and problem solving (McGuinness, 1986). Properties of the human cognitive and perceptual apparatus presumably also are important (e.g., Winn, 1994). For example, when applied to a particular problem, representations that are well organized and easy to scan presumably will be preferred to those that are disorganized and difficult to follow. In some instances, this criterion may override the representation choice specified by strict consideration of the problem's structure, as in the case of our nurses scenario discussed above. We suspect, however, that such overriding is the exception rather than the rule. For example, it only occurred for 1 of our 18 scenarios. Moreover, we predict that substitution of a structurally less appropriate representation for a more appropriate one will only facilitate problem solving if the substituted representation provides a substantially good fit to the structure of the problem in question, as is the case for our nurses scenario. In our view, problem structure provides the fundamental constraint for selecting a useful diagrammatic representation.

Issues On the Horizon

In closing, we briefly mention two additional directions for future research that we believe are important. In validating our structural analysis of the three spatial diagrams, it is important to use situations for which the structural properties uniquely specify a single representation. That is, one must begin by specifying the representation prototypes. It seems reasonable to presume, however, that some (perhaps many) situations that people encounter are ambiguous because applicability conditions for more than one representation are available in the input. Once the prototypes have been clearly specified, it will be important to determine how people resolve the inconsistencies present in such conflict situations. Finally, we should note that selecting an appropriate representation is the beginning of the problem-solving process, not the end (Novick, 2001). After we have a firm grasp of students' representational knowledge, it will be important to undertake studies of how the selected representation is constructed and used in the service of comprehension, reasoning, and problem solving.

APPENDIX A

Coding Guide for Representation Justification Responses (Structural Codes Only)

Each bracketed unit in a subject's protocol received three codes, e.g., M/Global/M. The first code indicated the representation to which the subject was referring: M (matrix), N (network), or H (hierarchy). The second code indicated the property to which the subject was referring. The third code indicated the value on that property to which the subject was referring. Typically, Code 3 was a particular representation. In a few cases, however, it was something else. In these cases, the alternate codes are indicated below. If the subject said exactly the opposite of what was indicated in the coding table (i.e., the subject negated something in the coding table), this was coded by using the most appropriate Code 3 code and putting a bar above it to indicate negation (e.g., $\overline{M}$ and $\overline{d}$). Negated codes were only allowed for the global structure, number of sets, and link type properties (see Appendix B). The possible values of Code 2 are printed in bold below. Under each of these codes are the possible associated values of Code 3 (underlined).

Global Structure **(Global)**

M: All values of one variable have the values of another variable in common; i.e., the representation expresses a factorial combination. For example, subjects might say things like ‘‘any of set X can be used with any of set Y’’ or ‘‘you want to see the combinations of the X things with the Y things.’’ The subject must clearly indicate that multiple things in one group are matched up with multiple things in another group.

N: The representation does not have any predefined, coherent structure. For example, subjects will say that the representation either (a) does not follow a coherent pattern, can take on any shape, is unstructured, or is random; or (b) has lines/arrows going all over the place or in all different directions. Do not use this code if the subject says something weaker like the representation is disorganized or complicated or has lots of arrows.

H: The representation is organized into levels beginning with a single root node (usually located at the top or right) that branches out to subsequent levels such that the identities of objects at one level depend on the identities of objects at the preceding level. Subjects must talk about the global structure of the representation in terms of dependencies or levels; reference to a single starting point by itself is not sufficient evidence for this code.

Building Block **(BBlock)**

M: A single cell/box denoting the intersection or combination of value *i* on one variable and value *j* on the other variable. Subjects must talk about the intersection point of two specific values, e.g., ‘‘look down the column for one value, across the row for another, and where they meet gives the answer.’’ Also, subjects must focus on only a single cell/box. If subjects talk about boxes/cells in general, consider whether LType/a is appropriate; do not use BBlock.

N: Two values/items and a (directional or nondirectional) relation between them. Subjects have to talk about two specific individual nodes linked together.

H: One value at a higher level, which gives rise to ≥2 values at a lower level; or ≥2 values at a lower level, which are narrowed down to one at a higher level; but not both. That is, three items and two directional relations arranged as an upside down or sideways ‘‘V.’’ Subjects have to talk about a specific node branching to multiple nodes, i.e., have to invoke an image of a V.

Number of Sets **(NSet)**

M: Two distinct sets of objects/concepts/items are being compared/combined. For example, the subject talks about ‘‘comparing two different [sets of] things.’’

N: All objects/concepts/items belong to a single set/group.

H: (a) All objects/concepts/items belong to a single set or (b) different levels of the hierarchy represent different sets of items.

Item/Link Constraints **(ILCon)**

M: Items within a set (i.e., within a row or column) can not be linked.

N: Any item may be linked to any other item (i.e., there are no constraints).

H: (a) Items at the same level cannot be linked, (b) items in nonadjacent levels cannot be linked, and (c) items in different ‘‘parts’’ of the tree (e.g., left vs. right side) cannot be linked.

Link Type **(LType)**

This code focuses on local connections between items.

a: There is a purely associative (i.e., nondirectional) link between items. Subjects must talk about a particular property being associated with two different things or about two different things being related in a specific way. Moreover, subjects must talk about this type of relation existing throughout the representation. If subjects focus on the relation between just two to three specific items, BBlock may be a more appropriate code.

d: There is a directional link between items. (Note that subjects often use the term ‘‘path’’ to refer to the link connecting two individual items.) References to ordering and use of directional words like ‘‘forward’’ and ‘‘from. . .to’’ are good clues that this code might be appropriate. ‘‘Arrow’’ also is a good key word, but only if other evidence also is available. Use ‘‘LType/b’’ if the subject talks about bidirectional links.

b: There is a bidirectional link between items. (Note that subjects often use the term ‘‘path’’ to refer to the link connecting two individual items.) The subject must specifically mention that a single link/line/path goes in both directions (e.g., from A to B and from B to A). If there is not clear evidence for this, use the more conservative ‘‘LType/d’’ code instead.

Absence of a Relation **(AbsRel)**

M: The absence of a relation between two items in different sets typically is indicated explicitly in the representation by a special mark (e.g., an ‘‘X’’) in the relevant cell. For example, ‘‘put an X if they don’t go together’’ or ‘‘to see if it doesn’t work, just look down the columns/rows.’’

N: The absence of a relation must be computed in all cases because there are no constraints on which items can be linked and there is no directly visible cue denoting the absence of a relation.

H: The absence of a relation generally is indicated implicitly due to the constraints on which items can be linked (but must be computed for non-linked items in adjacent levels).

Linking Relations **(LRel)**

This property focuses on a single node and the links going into and out of it.

b: Any number of lines/relations can enter and leave each node. That is, the relations involving a particular concept can be both one to many and many to one. (Because the number of lines into the node is what distinguishes LRel/b from LRel/o, at least as our subjects talk about it, a statement about the input links alone is sufficient evidence to give either LRel code.)

o: Only a single line can enter each node, but multiple lines can leave each node (i.e., all relations are one to many), or vice versa (all relations are many to one, as in elimination), but not both.

Path **(Path)**

This property focuses on the representation as a whole (unlike LType and LRel).

~p: This type of diagram does not allow one to represent or trace/follow an extended route connecting subsets of (>2) items.

p: This type of diagram allows one to represent or trace/follow an extended route connecting subsets of (>2) items. The word ‘‘path’’ by itself should not be used as evidence for this code because subjects often use ‘‘path’’ to refer to the link between two individual items. Note that explanations that implicitly include Path, like TravR, should not get a Path code unless a separate Path justification also is given.

Traversal **(TravR)**

This is a global property of the representation.

M: It does not really make sense to talk about traversing this type of representation.

N: Multiple paths from one node to another are possible; i.e., closed loops are allowed.

H: For any pair of nodes, A and B, there is only one path to get from A to B, i.e., there are no closed loops. For example, any mention of not being able to ‘‘loop around’’ or ‘‘can’t come back to the same item/starting node.’’

APPENDIX B

Details on How the Protocols Were Coded

We need to note that in addition to the 23 computer science majors and math educators, we also collected data from 12 more typical college students.

The data from this latter group of subjects are presented in Novick (2001). The protocol coding scheme, which is described in detail here, was developed using the full data set. In the body of this article, we only report coding information pertaining to the computer science and math education students. In discussing the details of the coding, however, it is difficult to separate the different subject groups because the protocols from all 35 subjects were coded together.

Developing the Coding Scheme and Coding the Protocols

Fifteen protocols (five from each subject population) were randomly selected to use in developing the coding scheme. Because we followed an iterative process to develop the coding scheme, it was possible that codes agreed upon at an early iteration were not consistent with changes made to the coding guide at later iterations. Therefore, after the coding guide was finalized, one coder reviewed the 15 ‘‘development’’ protocols to check that the codes conformed to the final version of the coding guide.

The coding criteria for each of the values of the 10 structural properties were, in most cases, simple elaborations of the property values given in Tables 1–3. For the global structure, number of sets, and link type properties, however, coders also were allowed to give ‘‘negated codes.’’ For example, a subject might say that a matrix would not be appropriate because the scenario has a levels or contingency structure. This type of response was coded as indicating that a matrix representation does not have the hierarchy value for the global structure property. Negated codes were not allowed for the other properties, primarily because subjects very rarely produced negations for them and we did not wish to unnecessarily proliferate coding categories.

Before a protocol was coded, there was a preliminary step in which each coder independently bracketed the statements for each scenario into what he/she thought were codable units. A codable unit was a clause, sentence, or set of consecutive sentences that the coder felt supported a single code. Then the two coders met to reconcile any discrepancies in the unit boundaries. It is important to note that discussion of where to begin and end codable units was made *without* talking about how those units might ultimately be coded.

After the unit boundaries for a protocol were finalized, the two coders independently coded the protocol without knowledge of the population from which the subject was recruited. Each bracketed unit received a three-part code. The first part indicated which representation the subject was talking about. The second part indicated the applicable coding category (e.g., global structure or link type). The third part indicated which value of the coding category the subject mentioned (e.g., matrix value or directional link). Sometimes subjects provided fairly explicit evidence for a property value by essentially paraphrasing what was in the coding guide in a fairly abstract manner.

More often, however, subjects conveyed the meaning indicated in the coding guide in language that was tied to the specific content of the scenario.

The Reliability of the Coding Scheme

After the coding scheme was finalized, the protocols for the remaining 20 subjects were coded by both coders in a predetermined random order. These data were used to compute the reliability of each coding category. As we mentioned in the body of the article, the item distinguishability property was rarely coded, and the two coders often did not agree on when its codes should be used. Of 360 justifications (20 subjects × 18 scenarios per subject), only 6 references to the item distinguishability property were coded. These 6 bracketed units were recoded before computing the reliabilities.

In computing the reliabilities, the two coders were said to agree if they gave exactly the same code for a particular subject and scenario. If the first two parts of the code were the same but the third part differed (e.g., N/LType/d vs. N/LType/a), we called this a partial agreement. There were only seven such pairs of codes across the 360 justifications in the reliability sample. For simplicity, these codes were excluded from the computation of reliabilities. That is, they were not counted as either agreements or disagreements. Overall, the two coders agreed on 73% of the structural codes given. This is quite good considering the number of codes and the complexity of the coding system: There were more than 50 possible structural codes that could be given because, for example, M/Global/M, M/Global/∼H, and H/Global/H are different codes. (The reliabilities were slightly lower for the data from the typical undergraduates.)

The one coding problem we encountered concerned the building block category, which was often confused with link type and (to a lesser extent) global structure. Although most of the specific discrepancies occurred fairly infrequently, for 17 scenarios Coder 1 coded M/LType/a and Coder 2 coded M/BBlock/M. Twelve of these 17 discrepancies came from computer science subjects, so it seemed possible that those subjects had a way of speaking that was potentially ambiguous. To determine whether Coder 1 was justified in interpreting the context in which the ‘‘ambiguous’’ statements were embedded as implying support for link type over building block, Coder 3, who has a masters degree in computer science, was trained in the distinction between these two codes for the matrix representation. This training was done using the justifications from the 15 subjects used by Coders 1 and 2 to learn the coding scheme. After the training, Coder 3 was given the 17 bracketed units that yielded discrepant codes for the 20 reliability subjects and was asked to categorize them as M/LType/a or M/BBlock/M. Coder 1 also recoded these units to ensure that the intended link type versus building block distinction was applied consistently. This resulted in Coder 1 changing two codes from link type to building block.

Coders 1 and 3 agreed on the coding of 14 of the 17 scenarios, or 82%. If Coder 3’s codes are substituted for Coder 2’s and the reliability of the building block category is recomputed, the new value for the computer science and math education subjects is .69 (up from .62), which is in line with the reliabilities for the other structural coding categories. Moreover, the fact that Coder 3 changed most of Coder 2's building block codes within this set to link type suggests that he would likely change many of Coder 2's remaining discrepant building block codes to something else, thereby likely further increasing the reliability of this coding category. Coder 3's responses also suggest that the disagreements probably reflect Coder 2's relative lack of familiarity with how computer science students talk rather than some fundamental problem with the coding scheme.

APPENDIX C

Examples of Coded Protocol Segments

Each protocol segment is divided into statements made by the subject (e.g., labeled 15.8 for Subject 15, Scenario 8) and statements made by the experimenter (labeled E). Usually, the experimenter only asked subjects why the representation they failed to choose was inferior for the scenario in question. In the transcripts, this is abbreviated as ''why X less good,'' where ''X'' is the first letter of the unchosen representation. The exact wording of the question is given under Method. A letter followed by a number indicates which representation the subject chose (Hierarchy, Matrix, or Network) and with what level of confidence (1–4).

As mentioned under Results, we coded both structural and nonstructural aspects of subjects' justifications. Several of the nonstructural codes appear in the protocols reprinted here, so we briefly describe all of those codes now. InstRep indicated that the subject instantiated the content of the scenario in a particular diagram (i.e., that the subject used the diagram as a representation). NotGood was used when the subject said the unchosen representation was not good but failed to give a codable reason for that statement. Easier was used when the subject simply said that the chosen representation would be easier to use. P/Exec/0 indicated evidence for what we termed executive monitoring, such as statements suggesting self-monitoring of one's justification, statements mentioning the goal of the scenario, and statements indicating an attempt to better understand the information presented. Xmpl indicated that the subject referred to some other specific situation (either an earlier scenario or a situation retrieved from memory) for which one of the spatial diagrams might be used. Mimic was used when subjects said that the unchosen representation would simply mimic the chosen representation. Finally, garb was our ''garbage'' code, indicating either that we could not figure out what the subject said or that what the subject said did not fit any of our coding categories. Unintelligible sections of the audiotapes are marked by ''{???}'' in the transcripts below.

Subject 15 (Math Education), Scenario 8 (see Table 6)

15.8: M3.

[because you're you're told that there's um two categories, A or B.][M/NSet/M] [and {???} I would put one on the side and one on the top.][M/InstRep/M] [if drugs from category A would be able to mix with category B you could be able to make note of that with maybe an X or a check.][M/LType/a] [whereas if they didn't you may want to leave it blank][M/AbsRel/M] [and you could actually see {???} how the pattern is formed. um which drugs are, it would be prescribed(?) to give.][P/Exec/0]

E: why N less good?

15.8: [because you're not, the organization is such that you can't exactly see any direct relationship between the two. I mean um, arrows are going to be go moving all over the place][N/Global/N] [and there's not really they're not really set up to show direct relationships between two uh between two things. they usually deal with three or more.][N/NSet/~M]

Subject 25 (Math Education), Scenario 10 (see Table 4)

25.10: M4.

[because again you're going to be comparing all the patients across, or possibly down, with all the doctors. and the result of the contact of each doctor with each patient][M/InstRep/M, M/Global/M] [is going to turn out with certain diagnoses which can go into the category.][M/LType/a] [and it's just organi, the information's going to be organized in such a way that that you can quickly reference which patient with which doctor. and if you wanted to see all the different ones for this patient, you could just look down the row and see which diagnoses. and if all the diagnoses were the same except for maybe this one, then you'd want to see which doctor was different from all the rest. and you can quickly come back and see his name.][M/Easier/M]

[it'd beat(?) the hierarchy or branching structure because the branching structure tends to lead and spread outward. but this this uh the diagram that they're asking for isn't something that's going to that you're going to want to show something spreading outward. because you're strictly comparing one group with another group. and the branching structure shows how one thing kind of leads to another possibly. but certainly isn't very good for for showing one group with another group. I don't even see, because each group will have certain number and this doesn't have any way to really compare that.][H/NSet/~M, H/LType/d]

Subject 8 (Computer Science), Scenario 11 (see Table 6)

8.11: N4.

[because um you have, back paths would be possible. . . . yeah because um because each um each child would have to have two parents in this case.][N/LRel/b]

[and um and a tree would get um would get long quite long for it.][H/NotGood/H] [and a digraph would represent it by um or trainers would have an out path going to their trainees, and trainees could have an in path coming from their trainers.][N/InstRep/N, N/LType/d]

E: why H less good?

8.11: [um because you got two possible parents for every um for every child in the tree. and um and it'd be quite hard to represent it in a tree logically.][H/LRel/o]

Subject 22 (Computer Science), Scenario 14 (see Table 6)

22.14: H4.

[first of all you have at the very top your Xantalin.][H/InstRep/H] [and then from there you would you'd break you'd break that part you'd break that drug down to like I don't know some some other breakdown of Xantalin. and you continue breaking that down until you get your cure or or whatever.][H/Global/H] [and if something is just going way left field, you could just pull that part out of the branch and then trace down another part of the branch to find to find a cure.][H/Path/p]

E: why N less good?

22.14: [OK. in the network or system of paths, you would just be just going nuts. first of all, after. you'd be tracing one path and it like leads to another.][N/Path/p]

[and and it with this type of system you don't want you don't want the breakdown to be coming back to a same node.][H/TravR/H] [therefore you want a different a different drug at each at each node.][H/InstRep/H]

[whereas a network or system of paths it ends up coming down to a node that can be traced from another node, which is not what you want.][N/LRel/b]

REFERENCES

Allwein, G., & Barwise, J. (Eds.). (1996). *Logical reasoning with diagrams.* New York: Oxford Univ. Press.

Bartram, D. J. (1980). Comprehending spatial information: The relative efficiency of different methods of presenting information about bus routes. *Journal of Applied Psychology,* **65,** 103–110.

Barwise, J., & Etchemendy, J. (1991). Visual information and valid reasoning. In W. Zimmermann & S. Cunningham (Eds.), *Visualization in teaching and learning mathematics* (pp. 9–24). Washington, DC: Mathematical Association of America.

Bertin, J. (1980). The basic test of the graph: A matrix theory of graph construction and cartography. In P. A. Kolers, M. E. Wrolstad, & H. Bouma (Eds.), *Processing of visible language* (Vol. 2, pp. 585–604). New York: Plenum. [Translated from the French by P. G. Allen.]

Bertin, J. (1981). *Graphics and graphic information-processing* (W. J. Berg & P. Scott, Trans.). Berlin: de Gruyter. (Original work published 1977)

Braxton, B., Gonsalves, P., Lipner, L., & Barber, J. (1995). *Math around the world: Teacher's guide.* Berkeley, CA: The Regents of the Univ. of California.

Broadbent, D. E., Cooper, P. J., & Broadbent, M. H. P. (1978). A comparison of hierarchical and matrix retrieval schemes in recall. *Journal of Experimental Psychology: Human Learning and Memory,* **4,** 486–497.

Butler, D. L. (1993). Graphics in psychology: Pictures, data, and especially concepts. *Behavior Research Methods, Instruments, & Computers,* **25,** 81–92.

Carroll, J. M., Thomas, J. C., & Malhotra, A. (1980). Presentation and representation in design problem-solving. *British Journal of Psychology,* **71,** 143–153.

Chartrand, G. (1985). *Introductory graph theory.* New York: Dover.

Cheng, P. C-H. (1996). Functional roles for the cognitive analysis of diagrams in problem solving. In G. W. Cottrell (Ed.), *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 207–212). Mahwah, NJ: Erlbaum.

Cormen, T. H., Leiserson, C. E., & Rivest, R. L. (1990). *Introduction to algorithms.* Cambridge, MA: MIT Press.

Day, R. S. (1988). Alternative representations. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 261–305). San Diego, CA: Academic Press.

Dufour-Janvier, B., Bednarz, N., & Belanger, M. (1987). Pedagogical considerations concerning the problem of representation. In C. Janvier (Ed.), *Problems of representation in the teaching and learning of mathematics* (pp. 109–122). Hillsdale, NJ: Erlbaum.

Elmasri, R., & Navathe, S. B. (1994). *Fundamentals of database systems* (2nd ed.). Menlo Park, CA: Addison–Wesley.

Epp, S. S. (1995). *Discrete mathematics with applications* (2nd ed.). Pacific Grove, CA: Brooks/Cole.

Ericsson, K. A., & Simon, H. A. (1993). *Protocol analysis: Verbal reports as data.* Cambridge, MA: MIT Press.

Glasgow, J., Narayanan, N. H., & Chandrasekaran, B. (Eds.). (1995). *Diagrammatic reasoning: Cognitive and computational perspectives.* Menlo Park, CA: AAAI Press/The MIT Press.

Guri-Rozenblit, S. (1988). The interrelations between diagrammatic representations and verbal explanations in learning from social science texts. *Instructional Science,* **17,** 219–234.

Hegarty, M., Carpenter, P. A., & Just, M. A. (1991). Diagrams in the comprehension of scientific texts. In R. Barr, M. L. Kamil, P. Mosenthal, & P. D. Pearson (Eds.), *Handbook of reading research* (Vol. 2, pp. 641–668). New York: Longman.

Holliday, W. G. (1976). Teaching verbal chains using flow diagrams and texts. *AV Communication Review,* **24,** 63–78.

Hurley, S. M., & Novick, L. R. (2001). Context and structure: The nature of students' knowledge about three spatial diagram representations. Manuscript submitted for publication.

Inhelder, B., & Piaget, J. (1964). *The early growth of logic in the child.* New York: Harper & Row. [Translated from the original French edition (1959) by E. A. Lunzer & D. Papert.]

Kindfield, A. C. H. (1993/1994). Biology diagrams: Tools to think with. *The Journal of the Learning Sciences,* **3,** 1–36.

Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science,* **11,** 65–99.

Lohse, G. L., Biolsi, K., Walker, N., & Rueter, H. H. (1994). A classification of visual representations. *Communications of the ACM,* **37**(12), 36–49.

Lynch, M. (1990). The externalized retina: Selection and mathematization in the visual documentation of objects in the life sciences. In M. Lynch & S. Woolgar (Eds.), *Representation in scientific practice* (pp. 153–186). Cambridge, MA: MIT Press.

Markman, A. B. (1999). *Knowledge representation.* Mahwah, NJ: Erlbaum.

McGuinness, C. (1986). Problem representation: The effects of spatial arrays. *Memory & Cognition,* **14,** 270–280.

Novick, L. R. (1990). Representational transfer in problem solving. *Psychological Science,* **1,** 128–132.

Novick, L. R. (2001). Spatial diagrams: Key instruments in the toolbox for thought. In D. L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 40, pp. 279–325). San Diego: Academic Press.

Novick, L. R., & Hmelo, C. E. (1994). Transferring symbolic representations across nonisomorphic problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **20,** 1296–1321.

Novick, L. R., Hurley, S. M., & Francis, M. (1999). Evidence for abstract, schematic knowledge of three spatial diagram representations. *Memory & Cognition,* **27,** 288–308.

Novick, L. R., & Morse, D. L. (2000). Folding a fish, making a mushroom: The role of diagrams in executing assembly procedures. *Memory & Cognition,* **28,** 1242–1256.

Polya, G. (1957). *How to solve it* (2nd ed.). Princeton, NJ: Princeton Univ. Press.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology,* **8,** 382–439.

Sanfey, A., & Hastie, R. (1998). Does evidence presentation format affect judgment? An experimental evaluation of displays of data for judgments. *Psychological Science,* **9,** 99–103.

Scanlon, D. A. (1989). Structured flowcharts outperform pseudocode: An experimental comparison. *IEEE Software,* **6**(Sept.), 28–36.

Schwartz, D. L. (1993). The construction and analogical transfer of symbolic visualizations. *Journal of Research in Science Teaching,* **30,** 1309–1325.

Schwartz, S. H. (1971). Modes of representation and problem solving: Well evolved is half solved. *Journal of Experimental Psychology,* **91,** 347–350.

Smith, C. D. (1987). Cartography in the prehistoric period in the old world: Europe, the Middle East, and North Africa. In J. B. Harley & D. Woodward (Eds.), *The history of cartography* (Vol. 1, pp. 54–101). Chicago: Univ. of Chicago Press.

Smith, D. E. (1925). *History of mathematics* (Vol. 2). Boston: Ginn and Co.

Smith, D. E., & Ginsburg, J. (1956). From numbers to numerals and from numerals to computation. In J. R. Newman (Ed.), *The world of mathematics* (Vol. 1, pp. 442–464). New York: Simon & Schuster.

Tversky, B. (2001). Spatial schemas in depictions. In M. Gattis (Ed.), *Spatial schemas and abstract thought* (pp. 79–111). Cambridge, MA: MIT Press.

Twyman, M. (1979). A schema for the study of graphic language (tutorial paper). In P. A. Kolers, M. E. Wrolstad, & H. Bouma (Eds.), *Processing of visible language* (Vol. 1, pp. 117–150). New York: Plenum.

Vessey, I., & Weber, R. (1986). Structured tools and conditional logic: An empirical investigation. *Communications of the ACM,* **29**(1), 48–57.

Weiss, M. A. (1999). *Data structures & algorithm analysis in Java.* Reading, MA: Addison–Wesley.

Winn, W. (1989). The design and use of instructional graphics. In H. Mandl & J. R. Levin (Eds.), *Knowledge acquisition from text and pictures* (pp. 125–144). Amsterdam, The Netherlands: Elsevier.

Winn, W. (1994). Contributions of perceptual and cognitive processes to the comprehension of graphics. In W. Schnotz & R. W. Kulhavy (Eds.), *Comprehension of graphics* (pp. 3–27). Amsterdam, The Netherlands: Elsevier.

Zaslavsky, C. (1973). *Africa counts: Number and pattern in African culture.* Boston: Prindle, Weber, and Schmidt.

Zaslavsky, C. (1998). *Math games & activities from around the world.* Chicago: Chicago Review Press.

(Accepted August 1, 2000, published online February 20, 2001)