

Mapping Conceptual to Spatial Relations in Visual Reasoning

Merideth Gattis and Keith J. Holyoak
University of California, Los Angeles

In 3 experiments, the authors investigated the impact of goals and perceptual relations on graph interpretation when people evaluate functional dependencies between continuous variables. Participants made inferences about the relative rate of 2 continuous linear variables (altitude and temperature). The authors varied the assignments of variables to axes, the perceived cause-effect relation between the variables, and the causal status of the variable being queried. The most striking finding was that accuracy was greater when the slope-mapping constraint was honored, which requires that the variable being queried be assigned to the vertical axis, so that steeper lines map to faster changes in the queried variable. The authors propose that graphs provide external instantiations of intermediate mental representations, enabling people to move from visuospatial representations to abstractions through the use of natural mappings between perceptual and conceptual relations.

Human reasoning appears to be based on multiple systems for representing information. Some types of inferences about spatial relationships appear to involve the manipulation of external or internal visuospatial displays (e.g., Hegarty & Just, 1993; Pinker & Finke, 1980; Shepard & Feng, 1972; Shepard & Metzler, 1971). In more abstract terms, people are also able to use spatial relations to reason about nonspatial relations. Such external representational media as graphs, diagrams, and schematic pictures are based on mappings between nonspatial and spatial relations, which allow inferences about the former to be drawn by means of perceptual operations performed on the latter. We use the term *visual reasoning* to refer to the use of visuospatial relations in making inferences about corresponding conceptual relations. In this study, we examined a particular example of visual reasoning: the use of graphical representations to evaluate functional dependencies between continuous variables. By investigating reasoning with graphs, we sought to understand how systematic concept-to-percept mappings underlie visual reasoning. We first briefly review what is known about general properties of visual reasoning and then discuss some particular issues that arise when conceptual relations are mapped to spatial relations in graphs.

Merideth Gattis and Keith J. Holyoak, Department of Psychology, University of California, Los Angeles.

This research was supported by National Science Foundation Grants USE-9150811 and SBR-9310614. Preliminary reports of these experiments were presented at the 34th Annual Meeting of the Psychonomics Society, Washington, DC, November 1993, and at the Sixteenth Annual Conference of the Cognitive Science Society, Atlanta, Georgia, August 1994. We thank David Shpall for his valuable technical assistance. Malcom Bauer and Barbara Tversky provided valuable comments on a draft of this article.

Correspondence concerning this article should be addressed to Merideth Gattis, who is now at Max Planck Institute for Psychological Research, Leopoldstr. 24, 80802 München, Germany. Electronic mail may be sent via Internet to gattis@mpipf-muenchen.mpg.de.

Cognitive Properties of Diagrams and Graphs

Visual reasoning is characterized by the use of spatial relations to highlight conceptual relations, to represent meaning, and to chunk information for computational efficiency. To explain why a diagram is sometimes more useful for reasoning than the equivalent textual representation, Larkin and Simon (1987) contrasted the informational and computational efficiency of sentential and diagrammatic representations. Representational equivalence, according to Larkin and Simon, is determined not only by whether the quantity of information in two representations is the same but also by whether that information is equally accessible and useful for the processes that will act on it. Informationally equivalent representations allow the same inferences to be made; computationally equivalent representations allow the same inferences to be made with the same time and effort.

Larkin and Simon (1987) contrasted the informational and computational efficiency of sentential and diagrammatic representations for a condition-action program exhibiting several computational requirements. These requirements included a search through an external representation of elements to identify those satisfying the conditions of stored production rules, recognition processes to match elements to conditions, and inference processes to execute the action portion of the rule. Larkin and Simon concluded that diagrammatic representations are more computationally efficient because they index information by location in a two-dimensional plane, in contrast to sequentially indexed sentential representations. In addition, relations that are implicit in sentential representations are often made explicit by diagrammatic representations, thus simplifying the search and recognition processes.

Investigations of visual reasoning have provided evidence that supports Larkin and Simon's (1987) explanation of efficiency differences between diagrammatic and sentential representations (Koedinger & Anderson, 1990). However, other work suggests that graphs and diagrams lead to differences in inferences that go beyond the recognition and search processes

on which Larkin and Simon focused. For example, Bauer and Johnson-Laird (1993) tested the efficiency of diagrams in supporting syllogistic reasoning. Participants were given either textual or diagrammatic representations of disjunctions and negations. The diagrams depicted circuits or similar abstract structures in which the premises were represented by physical constraints and validity was represented by a completed path. Such diagrams enable viewers to test potential conclusions based on alternative states of the world, or mental models, to see whether each possible conclusion is consistent with the constraints on the system. Participants who were given spatial analogs for disjunction, negation, and affirmation drew more accurate and rapid conclusions than did those who were given the textual syllogisms.

The differential performance observed by Bauer and Johnson-Laird (1993) suggests that such diagrams highlight the information used in reasoning with internal mental models and provide external support in managing the relations in working memory. Evidence for this conclusion was provided by a preliminary experiment in which Bauer and Johnson-Laird used a more arbitrary diagrammatic representation, which did not lead to superior performance for the diagram group relative to the text group. Bauer and Johnson-Laird attributed the difference in results obtained with the two types of diagrams to the fact that only the effective diagrams provided a nonarbitrary spatial analog for disjunction, and concluded that

The existence of the figural effect in the diagram condition strongly supports the notion that the figural effect is not due to linguistic form but is a more general phenomenon due to the order in which information enters working memory (1993, p. 377).

Glenberg and Langston (1992) reported additional evidence that diagrams play a special role in comprehension of text when the spatial relations in the diagram are analogous to important conceptual relations. These investigators constructed two types of diagrammatic representations of textual descriptions of procedures. One set of diagrams maintained the linear ordering inherent in the text by locating the steps one below the other, whereas the other set of diagrams highlighted the simultaneity of some steps by representing them as adjacent locations in the horizontal plane. After studying text only or text plus one set of diagrams, participants judged the validity of statements about the order of the steps in a procedure. Participants were more accurate at verifying the sequence of steps after studying a diagram only if they had studied the spatially analogous diagram; studying the linear diagram did not improve performance.

Glenberg and Langston (1992) argued that spatial diagrams improve text comprehension because they are external aids to the dynamic spatial representations, or mental models, derived from text and maintained in working memory. Readers construct mental models by mapping conceptual dimensions in the text onto spatial dimensions. The authors gave the example, "If the text describes particles that differ in energy, one (spatial) dimension can be used to represent energy, and the representational elements corresponding to the particles will be arrayed along that dimension" (Glenberg and Langston, 1992, p. 130). Such graph-like spatial arrays of concepts aid comprehension by improving management of working memory, by drawing attention to adjacent relations, and by strengthen-

ing connections between elements with similar values across some dimension.

The diagrammatic representations used in these studies of visual reasoning all exhibit many graph-like properties, most notably the mapping of concepts to spatial relations. Graphs provide a particularly pure case of visual reasoning because their form has evolved from natural cognitive constraints. Like pictures, graphs are simpler and more direct than other symbolic systems. By engaging built-in perceptual processing mechanisms, graphs function at an intermediate level of abstraction. Graphs exploit natural properties of the visual system in at least two ways. First, graphs promote computational efficiency for abstract problems by reducing the number of solution steps that sentential or mathematical procedures would require, or by reducing the number of dimensions or variables that must be explicitly compared. Utilizing the visual dimension to express information allows graphs to coordinate information on two or more dimensions. Integrating information across dimensions normally imposes a heavy cognitive load, but graphs reduce this load by integrating values on dimensions by means of points or lines that simultaneously represent a value or set of values on more than one dimension. Such visual integration, or chunking by gestalt grouping principles, allows people to reason about relations between two or more sets of data on those two or more dimensions.

Second, graphs also exploit vision by utilizing perceptual properties to convey conceptual relations. The "natural" mapping between the values on perceptual and conceptual dimensions is so automatic for some pairings, such as "greater area (more dots, thicker bars, or higher lines) equals greater quantity," that we often fail to recognize the relationship as a mapping rather than as an inherent equivalence of meaning. Other natural mappings involve more abstract pairings of perceptual and conceptual relations, such as "steeper equals faster," which is the focus of the present experiments. Mappings between perceptual and conceptual relations allow graphs to support reasoning about abstract concepts that are difficult to grasp directly, such as slope and function. For example, Pinker (1990) described several experiments by himself and Simcox (1983) indicating that global trends are identified by global features readily encoded by the visual system, whereas local values are identified by local features.

Although some have argued that graphs are often treated as pictures in educational settings (Clement, 1989; Leinhardt, Zaslavsky, & Stein, 1990), several studies indicate that graphs differ from pictures in important ways. Although graphs present information spatially, they are more abstract than simple pictorial representations. For example, the graph-like diagrams used by Glenberg and Langston (1992) involved mapping a property to a spatial dimension (specifically, mapping sequence or time to vertical and horizontal space). In more general terms, graphs are able to omit nonessential details and to highlight central higher order relations, such as rate of change in a variable, on the basis of a mapping from visuospatial properties of the graph to symbolically interpreted concepts.

Tversky and Schiano (1989) and Schiano and Tversky (1992) found evidence that the rules used for encoding graphs are not simply properties of the visual system but rather result from interpreting a figure as a graph. They demonstrated that the

label given to a diagrammatic representation invokes a particular set of rules or biases for encoding. Calling the representation a graph created a particular perceptual bias (the interpretation of a line orientation was biased toward 45°), whereas calling it a map or a figure either did not create any bias or biased the interpretation away from all orienting angles (including 45°). These results indicate that graphs constitute a symbol system that is associated with a specialized set of rules connecting the conceptual to the perceptual. In more general terms, we would argue, visual reasoning depends on systematic mappings between conceptual and spatial relations.

Conflicting Constraints on Graphical Representations

In this study, we contrasted pictorial and higher order spatial relations in graphs to show how graphs connect percepts to meaning. The issues we addressed concerned how people use graphical representations to evaluate functional dependencies between continuous variables. We placed different constraints on graphical representations in conflict so as to determine which constraints play the most important role in graphical reasoning. In particular, we pitted low-level pictorial correspondences between the domain being represented and the visuospatial display against the need to capture systematic correspondences between higher order relations in the represented domain and visuospatial relations in the display.

Figure 1 illustrates this conflict with a graph, typical of those used in teaching atmospheric science, representing functional dependencies between altitude and temperature. Textbooks and teachers invariably plot such altitude–temperature graphs with altitude on the y axis (e.g., Ahrens, 1988) because this preserves the low-level, pictorial correspondence between “up” in the world and “up” in the graph.

However, in the case of the functional dependency between altitude (the causal variable) and temperature (the effect), when the verticality of altitude is preserved, the rate of change in the effect (dependent variable) is the inverse of the slope of the data line. To be more general, if the independent variable (IV) is mapped to the x axis, and the dependent variable (DV) to the y axis, then the rate of change in the DV with respect to

the IV maps to the rate of change in y with respect to x (i.e., the slope of the line). We call this assignment the *slope-mapping constraint* because it ensures that judgments about rates can be based on the visually transparent mapping *steeper = faster*. Assigning the dependent variable to the y axis is a graphing convention extending back to the first known graphs. In 1765, Lambert graphed the evaporation of water as a function of temperature and explicitly noted that in such a graph the rate of evaporation could be taken from the slope of the line (Tuft, 1983). This pattern of axis assignment is seldom disregarded except in the case of horizontal bar graphs, where the assignment is reversed (Kosslyn, 1994). However, the standard pictorial-based assignment of altitude (the IV) to the vertical axis violates the slope-mapping constraint (because the rate of change in the DV will equal the reciprocal of the slope, instead of the slope).

In these experiments, we investigated visual reasoning about relative rates when the pictorial assignment of altitude to the vertical axis conflicted with the slope-mapping constraint. Our basic goal was to determine whether rate judgments are more accurate when the mapping from conceptual to perceptual variables preserves either (a) low-level pictorial correspondences or (b) the more abstract correspondence dictated by the slope-mapping constraint. In addition, the experiments were designed to distinguish among several factors that might underlie the slope-mapping constraint. The most general version of the constraint would imply that the variable for which the rate is being queried (e.g., temperature in the above example) be assigned to the y axis. However, in addition to serving as the queried variable, temperature plays the role of the effect in a cause–effect relationship. Not all functional relationships between continuous variables involve cause–effect relations. Even for those that do, it is possible for people to make not only predictive inferences from causes to their effects but also diagnostic inferences from effects to their causes (Waldmann & Holyoak, 1992). For example, one could sensibly query how rate of change in altitude (the cause) varies with change in temperature (the effect). Experiment 2 com-

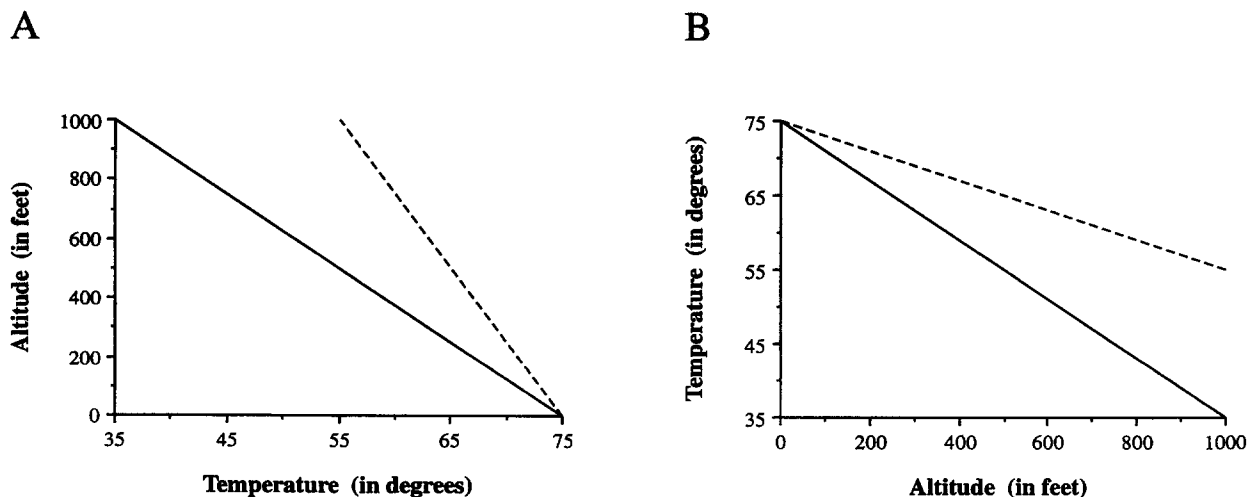


Figure 1. Graphs that violate (A) or conform with (B) the slope-mapping constraint.

pared rate judgments for causal versus noncausal functional relationships, and Experiment 3 compared rate judgments for causal relationships when the effect or else the cause was queried.

Experiment 1: Conflicting Constraints

Method

Participants. Sixty-six undergraduate psychology students at the University of California, Los Angeles (UCLA) participated in exchange for course credit. Approximately half of the participants were science majors and half were nonscience majors, distributed randomly across conditions.

Materials, procedure, and design. Participants were given graphs representing the functional dependency between atmospheric altitude and air temperature, along with a brief explanation of their relationship. Because this dependency is a loose causal relation in that altitude does not directly cause temperature, the causal relation between altitude and temperature was implied rather than made explicit in these instructions (see the Appendix for the complete instructions). In fact a complex interaction of several factors, including altitude, determines air temperature. In the case of a moving air parcel, the change in altitude causes a change in temperature. The graphs either represented altitude on the y axis in accordance with the atmospheric-science tradition of preserving verticality, but in violation of the slope-mapping constraint (Figure 1A), or on the x axis, in accordance with the graphing convention (following from the slope-mapping constraint) of mapping the IV to the x axis (Figure 1B). Because of the causal relation between the variables represented in these graphs, altitude was the IV and temperature was the DV.

On the first page presented to participants, a single, solid data line was drawn across the graph, representing a single value of y for each value of x . The line displayed an inverse linear relationship between altitude and temperature, with higher temperatures at lower altitudes and lower temperatures at higher altitudes.¹ All participants were first asked to identify a particular point on the graph and then to make a judgment about the direction of change for one of the variables represented in the graph. These questions were used to screen and eliminate participants who could not identify the component information required for rate judgments, so as to maximize the likelihood that any performance differences observed in the experiment would be attributable to our manipulation, rather than to variations in participants' basic graphing skills.

On the second page, participants were shown another graph and asked to perform one of two rate-judgment tasks. In the drawing task, participants were given the original graph with the single, solid data line accompanied by a description of a second data set in which temperature changed more slowly relative to altitude, and they were asked to draw a line that could represent such a data set. In the multiple-choice task, participants were given the original graph with the solid data line and an additional, dotted second line and were asked to judge whether the second set of data (the dotted line) represented a slower or faster rate of change of temperature than the first.

Participants were assigned in approximately equal numbers to four conditions defined by the factorial combination of axis assignment (altitude on x or else y axis) and task (drawing or multiple choice). Assigning altitude (the IV) to the y axis preserved the pictorial correspondence of *up = up* at the cost of violating the slope-mapping constraint of *steeper = faster*, whereas assigning altitude to the x axis reversed these properties. We predicted that judging rate of change would be more dependent on the higher order mapping than on the pictorial correspondence, so that accuracy would be higher when altitude was on the x axis.

Results and Discussion

Data from 8 participants were excluded on the basis of the screening questions (4 participants in the draw condition with altitude on y and 2 in each of the other three conditions). As predicted, participants were more accurate in both rate tasks when the IV (altitude) was assigned to the x axis and the DV (temperature) to the y axis. All analyses of solution frequencies were based on the G^2 statistic (maximum-likelihood chi-square). For the drawing task, the percentage of participants who drew a line showing the correct rate relation was 78% when altitude was on x and 33% when it was on y , $G^2(1, N = 18) = 5.60, p < .025$. Similarly, the percentages of correct choices in the multiple-choice task by subjects in the two conditions were 100% and 64%, respectively, $G^2(1, N = 30) = 6.86, p < .01$. Major (science and nonscience) had no effect on task performance. These results reveal that when the two pressures conflicted in a graphical representation of rate of change, preserving the pictorial relation of verticality was less important than preserving the more abstract correspondence between steepness of lines and speed of rate change.

Experiment 2: Testing Alternative Interpretations

Experiment 2 was designed to distinguish the slope-mapping constraint from other factors with which it was correlated in Experiment 1. In the cover story used in Experiment 1, as in many naturally occurring situations, the IV (altitude) was conceptually the cause of the DV (temperature), which was thus interpreted as the effect. According to the slope-mapping constraint, the crucial variable determining the ease of answering questions about relative rate is that the IV must be assigned to the x axis and the DV (the variable queried) to the y axis. However, the cover story used in Experiment 1 created a natural confounding between the assignment of the IV to the x axis with (a) the assignment of a specific dimension, altitude, and (b) the assignment of the variable interpreted as a cause. Experiment 2 was performed to disentangle these alternative conceptualizations of the factors underlying the axis-assignment preference observed in Experiment 1. By manipulating cover stories, we examined whether the results depended on the existence of a conceptual relation of cause and effect between the IV and DV and whether there is a preference for a particular causal directionality in making judgments of relative rate.

Method

Participants. One hundred seventy students in undergraduate psychology courses at UCLA participated in exchange for course credit.

Materials, procedure, and design. The two graphs used in Experiment 1 were also used in Experiment 2. In the causal condition the x and y axes were labeled in the same way as in Experiment 1. In the noncausal condition the x and y axes were labeled simply as A and B . To manipulate the causal direction between altitude and temperature, two cover stories were created for the causal condition. One was the

¹ An inverse linear relationship was used because pilot studies revealed that performance of adult participants was at ceiling for rate judgments of direct relationships.

version used in Experiment 1, in which altitude loosely causes temperature. A second version described the movement of a hot air balloon, a scenario in which atmospheric air temperature (relative to a constant balloon air temperature) causes changes in altitude (of the balloon). The functional relations described in the latter story were nearly identical to those described in the former, except that the causal direction was reversed. In the case of a hot air balloon, change in the air temperature loosely causes change in the altitude of the balloon. No cover story was used for the noncausal conditions.

Participants were assigned to one of eight conditions defined by the factorial combination of axis queried (x vs. y), causal status of the cover story (causal or noncausal), and the variable queried (altitude or temperature for the causal conditions and A or B for the noncausal conditions). Participants in four conditions received one of the two causal cover stories (making either altitude or temperature the cause). In each case, the variable assigned to the x axis was varied. Participants in the other four conditions did not receive any causal cover story, so the graph simply represented a functional dependency. Participants in these noncausal conditions were given graphs identical to those in the four causal conditions but without the cover story or the altitude and temperature labels. The x and y axes were labeled A and B, with the label assignments counterbalanced across each condition.

The two graph interpretation questions used in Experiment 1 were presented on the first page, along with a graph of a single set of data. All subjects were presented a second page with that same graph together with an additional set of data represented as a dotted line. As in the multiple-choice task used in Experiment 1, participants were instructed to choose the correct rate comparison description for the second data line from two alternatives (faster or slower than the first line). This question always probed the rate of change in the DV (the effect in the causal conditions) with respect to the IV. In the noncausal conditions, the DV was defined as the variable (A or B) being queried. The more difficult drawing task that had been used in Experiment 1 was not used.

The design was such that among the four causal conditions, two had the cause (IV) assigned to the x axis and the effect (DV) assigned to the y axis, in accord with the slope-mapping constraint, and two had the reverse assignment. When altitude was the cause, the slope-mapping constraint conflicted with the pictorial correspondence of $up = up$, as was the case in Experiment 1, whereas when altitude was the effect the two pressures converged. The design of Experiment 2 thus enabled us to assess whether preserving the pictorial correspondence might at least add some additional benefit when it was in accord with the slope-mapping constraint. Among the four matched noncausal conditions, two satisfied the slope-mapping constraint (queried variable on the y axis) and two violated it.

Results and Discussion

Data from 34 participants were excluded on the basis of the screening questions (9 participants in the causal conditions, 5 of whom were queried on temperature with altitude on y , and 25 in the noncausal conditions who were distributed evenly across the four cells). Table 1 presents the mean percentage correct on the rate-judgment task for each of the conditions in Experiment 2. As in Experiment 1, participants made rate judgments more accurately (93% vs. 46%) when the slope-mapping constraint was satisfied (DV, or queried variable, on the y axis) than when it was not, $G^2(1, N = 136) = 39.6, p < .0001$. This effect held regardless of whether a causal cover story was provided and whether altitude or temperature was the cause. No other main effects or interactions approached significance, $G^2(1) < 1$, in each case.

Table 1
Percentage Correct Rate Judgments in Each Condition of Experiment 2

Causal relation and axis assignment	% correct
Causal status	
Temp. is effect (queried)	
Altitude on x	85 ^a
Altitude on y	50
Altitude is effect (queried)	
Altitude on x	35
Altitude on y	95 ^a
Noncausal status	
A queried	
A on x	94 ^a
A on y	43
B queried	
A on x	57
A on y	100 ^a

Note. A and B refer to noncausal variable labels. Temp. = temperature.

^aThis percentage is for a condition in which the slope-mapping constraint was satisfied.

These results thus confirm the central finding of Experiment 1: Judgments of relative rate were more accurate when the axis assignments honored the slope-mapping constraint. In addition, the results of Experiment 2 indicate that preserving the pictorial mapping of altitude to the vertical axis did not convey any significant added benefit, either when the slope-mapping constraint was satisfied or when it was violated. Moreover, the results rule out two alternative interpretations of the results obtained in Experiment 1. The crucial factor is that the IV must be assigned to the x axis, in accord with the slope-mapping constraint; it is not crucial that altitude per se or a variable interpreted as a cause be so assigned. Thus in the noncausal conditions, as in the causal conditions, participants were more accurate in judging relative rate when the IV was assigned to the x axis and the DV (the queried variable) was assigned to the y axis.

Experiment 3: Querying Cause Versus Effect

Although the results of Experiment 2 did not reveal any influence of a causal cover story on rate judgments, it would be premature to rule out any role of causal relations in graph interpretation. Prior experience with graphs in which the cause variable was always placed on the x axis could have unconsciously carried over to graphs without any explicit causal interpretation, so that participants may have assumed that the IV in the noncausal condition was in fact a cause. In addition, Experiments 1 and 2 only tested participants' ability to make inferences about the rate of change in an effect as its cause varied. This type of question involves a predictive inference from a cause to its effect. However, people are also capable of making diagnostic inferences, judging variations in a cause by variations in its effect (Waldmann & Holyoak, 1992). Assuming that people tend to code causal knowledge in the cause-to-effect direction, diagnostic inferences that depend on back-

ward reasoning from effect to cause may be more difficult, especially when cast as judgments about higher order relations, such as rate of change. It is therefore possible that people will be especially sensitive to violations of graphing conventions in the context of judging the rate at which a cause changes with its effect. Accordingly, in Experiment 3 we asked participants to reason about a change in cause as a function of a change in effect, as well as the opposite case. Experiment 3 was conducted to examine the combined influence of three factors on graph interpretation: the slope-mapping constraint, preserving pictorial correspondences with the real world, and a possible asymmetry in causal inference.

Method

Participants. One hundred fifty-nine students in undergraduate psychology courses at UCLA participated in Experiment 3 in exchange for course credit.

Materials, procedure, and design. The four causal conditions from Experiment 2 were repeated in this experiment, with the same graphs and axis labels. The two cover stories used in Experiment 2 to manipulate the causal direction between altitude and temperature were again used in Experiment 3, with slight changes to create a context in which participants would find it plausible to be asked about the change in the cause as a function of the change in the effect. The multiple-choice task used in the previous experiments was used once again in Experiment 3.

Participants were assigned to one of eight conditions defined by the factorial combination of the axis queried (x vs. y), the causal status of the queried variable (cause or effect), and the variable interpreted as the effect (altitude or temperature). Experiment 3 was identical to Experiment 2 in all other respects.

Results and Discussion

Data from 20 participants were excluded on the basis of the screening questions (6 participants in the condition in which altitude was on x and queried as the effect, 5 in the condition in which altitude was on x and temperature was queried as the cause, and 0 to 2 participants in each of the remaining six conditions). Table 2 presents the percentages of correct rate judgments for each of the eight conditions. As in the previous experiments, accuracy was higher (92% vs. 54%) when the queried variable was assigned to the y axis rather than to the x axis, $G^2(1, N = 139) = 26.2, p < .0001$. Overall, this effect did not vary significantly with the identity or causal status of the queried variable, $G^2(1, N = 139) = 1.90, p > .1$, and $G^2(1, N = 139) = 3.09, p < .10$, respectively. No other main effects or interactions achieved statistical significance, although there was a tendency toward especially low accuracy when altitude was queried as a cause assigned to the x axis, $G^2(1, N = 139) = 2.91, p < .10$. Overall, querying the cause did not significantly reduce accuracy relative to querying the effect, $G^2(1, N = 139) = 2.78, p < .10$.² The interpretation of relative rate in a graphical representation did not appear to depend on whether the question required predictive reasoning (from cause to effect) or diagnostic reasoning (from effect to cause).

The results of Experiment 3 thus support the most general version of the slope-mapping constraint. If one follows the conventional interpretation of the cause as the IV and the effect as the DV, our findings demonstrate that it is not

Table 2
Percentage Correct Rate Judgments in Each Condition of Experiment 3

Causal relation and axis assignment	% correct
Rate comparison queries effect	
Temp. is effect	
Altitude on x	90 ^a
Altitude on y	61
Altitude is effect	
Altitude on x	64
Altitude on y	89 ^a
Rate comparison queries cause	
Temp. is cause	
Altitude on x	93 ^a
Altitude on y	72
Altitude is cause	
Altitude on x	22
Altitude on y	95 ^a

Note. Temp. = temperature.

^aThis percentage is for a condition in which the slope-mapping constraint was satisfied.

essential for the IV to be assigned to the x axis and the DV to the y axis. Rather, the crucial requirement for satisfying the slope-mapping constraint is that the variable being queried (whether the DV or IV, as defined by the causal interpretation) be assigned to the y axis.

General Discussion

The results of this study help to illuminate the component processes in graph interpretation and the manner in which visual reasoning relies on mappings from conceptual to spatial relations. We found that the ease of drawing inferences about rate of change depended on the nature of the mapping between a conceptual dimension (rate of change in a continu-

² The nonsignificant trends involving the causal status of the queried variable (both as a main effect and as a component of interactions) reflect the especially low accuracy in the condition in which altitude was queried as a cause assigned to the x axis (22% correct; see Table 2). To assess the reliability of this finding, an additional experiment with different participants was performed. The four conditions of Experiment 3 in which the cause was queried (bottom half of Table 2) were replicated. A total of 73 participants served in the experiment, with from 16 to 20 participants assigned to each condition. Data from 7 participants were excluded on the basis of the screening questions (4 participants in the condition in which altitude was on x and temperature was queried as the cause). The percentage correct rate judgments in these conditions when temperature was the cause were 94% when altitude was on the x axis and 53% when altitude was on the y axis; when altitude was the cause, the corresponding figures were 50% and 75%. These results thus replicated the key finding of Experiment 3, with participants achieving greater accuracy when the slope-mapping constraint was satisfied than when it was not, $G^2(1, N = 66) = 10.32, p < .01$. Unlike in Experiment 3, there was no tendency for the condition in which altitude was queried as the cause assigned to the x axis to be especially difficult, $G^2(1, N = 66) < 1$. We therefore consider the especially low performance observed in this condition of Experiment 3 to be unreliable.

ous variable) and the visuospatial dimension of slope of a line. Furthermore, the optimal mapping was dependent on the reasoner's goal, as manipulated by varying which variable was queried with respect to its rate of change. Our most consistent and powerful finding was that accuracy was greater when the slope-mapping constraint was honored, which requires that the variable being queried—usually the effect or dependent variable, but potentially the cause instead—is assigned to the vertical axis, so that steeper lines map to faster changes in the queried variable. This constraint dominates when it conflicts with others, such as preserving the low-level mapping of altitude onto the vertical axis.

These results support the basic conclusion that external visual representations such as graphs, diagrams, and schematic pictures have special properties that, when used appropriately, can facilitate reasoning about abstract higher order relations. Such external representations mediate between internal representations at the visuospatial and symbolic levels by utilizing visuospatial relations to convey conceptual information. Sometimes the dependency of conceptual expression on perceptual properties can lead to confusion between graphs and pictures (Clement, 1989), but for the most part this relationship provides an efficient representational mechanism.

An important issue for future research is to determine whether the slope-mapping constraint observed here is based on a deep property of the cognitive system, or whether it develops as a product of experience with graphs and the conventions of axis assignment. Educational studies reporting that graph comprehension is a particularly troublesome task for students indicate that graphing conventions are not well learned. Leinhardt et al. (1990) reviewed many educational studies documenting students' poor performance on many graphing tasks, especially those that involve recognizing patterns or abstracting information from graphs. In contrast, Tversky, Kugelmass, and Winter (1991) have obtained evidence that graphing conventions for mapping increases are based on natural cognitive correspondences. Similarly, Lakoff and Johnson (1980) have argued that the metaphorical mapping between the vertical dimension of space and the upward direction of abstract dimensions is likely to be a cognitive universal (see also Miller & Johnson-Laird, 1976). From this perspective, our results show that when we must choose between maintaining this spatial mapping either at a pictorial level (*up = up*) or at the level of more abstract relations (*steeper = faster*), the latter mapping is more important for inferring rates. If this view is correct, we would expect that the nature of the human cognitive system has led to this graphing convention, which is nonarbitrary.

A developmental cross-linguistic study of performance on a graphic construction task provides some preliminary support for this view. Tversky et al. (1991) asked children and adults from three different language groups (Arabic, English, and Hebrew) to place stickers on a piece of paper to represent spatial, temporal, and quantitative relations among various entities. Their results revealed influences of language-specific directionality (left-to-right vs. right-to-left order of reading and writing) on sticker placement for temporal relations. However, what is more important in the present context is that

the study revealed a universal association of "more" or "better" with the upward direction for all language groups and ages, supporting the possibility that this mapping is deeply rooted in human cognition. More work is required to determine whether the dominance of the slope-mapping constraint over pictorially based mappings is also universal, or whether it is subject to developmental change. Further developmental studies of the relative ease of learning to interpret graphs under alternative conventions, such as those for assignment of axes, may cast light on the origins of natural rules for graph construction and interpretation.

This issue of the developmental origins of graph conventions is related to the larger question of whether other perceptual-conceptual linkages found in graphs arise from mental representations that mediate the pictorial and the symbolic. Graph theorists have fostered the belief that rules for graphing are the products of explicit design decisions by those schooled in the art and science of representation, imagery, or symbols (e.g., statisticians, graphical designers, and semioticians). It is possible, however, that in constructing graphs and adopting conventions for their use people have been tacitly enhancing natural representations that already existed in the human mind. Evaluating the latter possibility will require searching for other rules of graphing that may depend on analogical mappings between perceptual and conceptual relations. In addition to studies addressing the above issues, deeper understanding of the nature of visual reasoning will require investigations of the differences between processing of relatively direct visual displays, such as realistic pictures and drawings, figures, maps, and some diagrams, versus more abstract visual displays, such as schematized pictures and diagrams, graphs, and mental models. Evidence from such studies may reveal that visual reasoning is generally based on procedures for finding analogical mappings between conceptual relations and physical space.

References

- Ahrens, C. D. (1988). *Meteorology today: An introduction to weather, climate, and the environment* (3rd ed.). St. Paul, MN: West.
- Bauer, M. I., & Johnson-Laird, P. N. (1993). How diagrams can improve reasoning. *Psychological Science*, 6, 372-378.
- Clement, J. (1989). The concept of variation and misconceptions in cartesian graphing. *Focus on Learning Problems in Mathematics*, 11, 77-87.
- Glenberg, A. M., & Langston, W. E. (1992). Comprehension of illustrated text: Pictures help to build mental models. *Journal of Memory and Language*, 31, 129-151.
- Hegarty, M., & Just, M. A. (1993). Constructing mental models of machines from text and diagrams. *Journal of Memory and Language*, 32, 717-742.
- Koedinger, K. R., & Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science*, 14, 511-550.
- Kosslyn, S. M. (1994). *Elements of graph design*. New York: W. H. Freeman.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65-99.

- Leinhardt, G., Zaslavsky, O., & Stein, M. K. (1990). Functions, graphs, and graphing: Tasks, learning, and teaching. *Review of Educational Research, 60*, 1–64.
- Miller, G. A., & Johnson-Laird, P. N. (1976). *Language and perception*. Cambridge, MA: Harvard University Press.
- Pinker, S. (1990). A theory of graph comprehension. In R. Freedle (Ed.), *Artificial intelligence and the future of testing* (pp. 73–126). Hillsdale, NJ: Erlbaum.
- Pinker, S., & Finke, R. A. (1980). Emergent two-dimensional patterns in images rotated in depth. *Journal of Experimental Psychology: Human Perception and Performance, 6*, 244–264.
- Schiano, D. J., & Tversky, B. (1992). Structure and strategy in encoding simplified graphs. *Memory & Cognition, 20*, 12–20.
- Shepard, R. N., & Feng, C. (1972). A chronometric study of mental paper folding. *Cognitive Psychology, 3*, 228–243.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science, 171*, 701–703.
- Simcox, W. A. (1983). *A perceptual analysis of graphic information processing*. Unpublished doctoral dissertation, Tufts University, Medford, MA.
- Tufte, E. R. (1983). *The visual display of quantitative information*. Cheshire, CT: Graphics Press.
- Tversky, B., Kugelmass, S., & Winter, A. (1991). Cross-cultural and developmental trends in graphic productions. *Cognitive Psychology, 23*, 515–557.
- Tversky, B., & Schiano, D. J. (1989). Perceptual and cognitive factors in distortions in memory for graphs and maps. *Journal of Experimental Psychology: General, 118*, 387–398.
- Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General, 121*, 222–236.

Appendix

Cover Stories for Experiments 1–3

Experiment 1

The following graph is called a sounding. A sounding is collected from a vertical slice of air by taking measurements of air temperature at regular intervals from a balloon as it ascends through the atmosphere. It shows the atmospheric air temperature at different altitudes. The line plotted in this graph represents the change in temperature as the altitude changes. Atmospheric scientists use this data to determine how air parcels will behave in a particular environment.

This is the same graph that you just saw with additional data from another sounding. The data from the second sounding is represented by the dotted line. Circle the letter next to the statement that best describes that data. In the second sounding, the air cools:

- at a faster rate than the first sounding.
- at a slower rate than the first sounding.

Experiment 2

Temperature Causes Altitude Story

This is a graph of the movement of a hot air balloon through the atmosphere. It shows the balloon's altitude at different temperatures of the surrounding air. Remember that hot air balloons operate on the principle that hot air is less dense than cool air and will rise. Therefore a balloon filled with air that is hotter than the surrounding air will be lifted into the atmosphere. This means that a balloon's altitude in the atmosphere is an effect of the temperature difference between the air in the balloon and the surrounding air. The air in this balloon is at a constant temperature. The line plotted in this graph represents the relationship between altitude and temperature. We can use this data to determine how a particular balloon will behave in a particular environment.

It turns out that there is an additional factor besides the air temperature difference that also determines the rate of ascent and descent: the weight of the balloon. As weight increases, the rate of altitude change decreases. This is the same graph that you just saw with additional data from a second balloon sent up on the next day. The data from the second balloon is represented by the dotted line. Circle the letter next to the statement that best describes that line. In the second line, **altitude** changes:

- at a faster rate.
- at a slower rate.

Altitude Causes Temperature Story

This is a graph of air temperature taken from a balloon at regular intervals as it ascends through the atmosphere. It shows the atmospheric air temperature at different altitudes. Remember that what makes air warm is that sunlight warms the earth's surface and the surface, in turn, warms the air above it. Therefore air that is closer to the earth is warmer and air at higher points in the atmosphere is cooler. This means that air temperature is determined by altitude. The line plotted in this graph represents the relationship between temperature and altitude. We can use this data to determine how a particular air parcel will behave in a particular environment.

It turns out that there is an additional factor that determines the air temperature: the moisture in the air. As moisture increases, the rate of temperature change decreases. This is the same graph that you just saw with additional data from a second balloon sent up on the next day. The data from the second balloon is represented by the dotted line. Circle the letter next to the statement that best describes that data. In the second line, **temperature** changes:

- at a faster rate.
- at a slower rate.

Experiment 3

Temperature Causes Altitude Story

First graph. In your new job as lab assistant in the atmospheric sciences department, your job is to collect and interpret data. This is a graph of data from an experiment investigating how a hot air balloon moves through the atmosphere. It shows the balloon's altitude at different temperatures of the surrounding air. Remember that hot air balloons operate on the principle that hot air is less dense than cool air and will rise. Therefore a balloon filled with air that is hotter than the surrounding air will be lifted into the atmosphere. These data were collected while the air inside the balloon was held at a constant temperature. The plot shows how changes in the temperature of the outside air produced changes in the altitude of the balloon.

Second graph. This morning your professor has collected more data in a new experiment. He sent up another hot air balloon under new conditions in order to obtain more information about how changes in temperature cause changes in altitude of a balloon. Since he is a very busy guy, he isn't here to work with you. He just left this

graph, which plots the new data along with yesterday's data. The solid line represents the data collected in yesterday's experiment (which is the same as the graph on the previous page). The dotted line represents the new results from this morning's experiment. The professor wants you to interpret the results. In particular, he would like to know about the apparent differences between the results of the two experiments. You decide that one way of examining this difference is to look at how altitude varies. From the graph, you conclude that as temperature changes, **altitude** changes

- a. at a faster rate
- b. at a slower rate

in today's experiment (dotted line) compared to yesterday's experiment (solid line).

Altitude Causes Temperature Story

First graph. In your new job as lab assistant in the atmospheric sciences department, your job is to collect and interpret data. This is a graph of air temperature taken from a balloon at regular intervals as it moves through the atmosphere. It shows the atmospheric air temperature at different altitudes. Remember that what makes air warm is that sunlight warms the earth's surface and the surface, in turn, warms the air above it. Therefore air that is closer to the earth is warmer and air at higher points in the atmosphere is cooler. This means that air

temperature is determined by altitude. The plot shows how changes in the altitude produced changes in the temperature of the air.

Second graph. This morning your professor has collected more data in a new experiment. He sent up another hot air balloon under new conditions in order to obtain more information about how changes in altitude cause changes in temperature. Since he is a very busy guy, he isn't here to work with you. He just left this graph, which plots the new data along with yesterday's data. The solid line represents the data collected in yesterday's experiment (which is the same as the graph on the previous page). The dotted line represents the new results from this morning's experiment. The professor wants you to interpret the results. In particular, he would like to know about the apparent differences between the results of the two experiments.

You decide that one way of examining this difference is to look at how temperature varies. From the graph, you conclude that as altitude changes, **temperature** changes

- a. at a faster rate
- b. at a slower rate

in today's experiment (dotted line) compared to yesterday's experiment (solid line).

Received September 1, 1994

Revision received December 16, 1994

Accepted January 19, 1995 ■

Low Publication Prices for APA Members and Affiliates

Keeping you up-to-date. All APA Fellows, Members, Associates, and Student Affiliates receive—as part of their annual dues—subscriptions to the *American Psychologist* and *APA Monitor*. High School Teacher and International Affiliates receive subscriptions to the *APA Monitor*, and they may subscribe to the *American Psychologist* at a significantly reduced rate. In addition, all Members and Student Affiliates are eligible for savings of up to 60% (plus a journal credit) on all other APA journals, as well as significant discounts on subscriptions from cooperating societies and publishers (e.g., the American Association for Counseling and Development, Academic Press, and Human Sciences Press).

Essential resources. APA members and affiliates receive special rates for purchases of APA books, including the *Publication Manual of the American Psychological Association*, and on dozens of new topical books each year.

Other benefits of membership. Membership in APA also provides eligibility for competitive insurance plans, continuing education programs, reduced APA convention fees, and specialty divisions.

More information. Write to American Psychological Association, Membership Services, 750 First Street, NE, Washington, DC 20002-4242.