Cognitive Science 28 (2004) 589–610

COGNITIVE SCIENCE

http://www.elsevier.com/locate/cogsci

Mapping relational structure in spatial reasoning

Merideth Gattis*

Max Planck Institute for Psychological Research, University of Sheffield, Sheffield, UK

Received 9 August 2003; received in revised form 20 January 2004; accepted 9 February 2004

Available online 19 March 2004

Abstract

Three experiments investigated whether the similarity of relational structures influences the interpretation of spatial representations. Adults were shown diagrams of hand gestures paired with simple statements and asked to judge the meaning of new gestures. In Experiment 1 the gestures were paired with active declarative statements. In Experiment 2, the gestures were paired with conjunctive and disjunctive relations. Experiment 3 used statements similar to those used in Experiment 1, but eliminated the initial object-to-object mapping provided in Experiments 1 and 2. In all three experiments, most participants chose an interpretation that set up a parallel relational structure between the gesture and its meaning: spatial elements were paired with conceptual elements and spatial relations were paired with conceptual relations. These results are consistent with the hypothesis that similarity of relational structures influences spatial reasoning, and have implications for analogical reasoning, diagrammatic reasoning, and language processing.

Keywords: Spatial reasoning; Diagrammatic reasoning; Conjuctive and disjuctive reasoning; Iconicity; Relational structure; Structure mapping

Spatial reasoning is the use of spatial representations, such as diagrams, graphs, and gestures, to reason. Research on spatial reasoning has frequently asked how spatial representations improve reasoning. To answer that question, cognitive scientists have usually analyzed the computational processes involved in reasoning with diagrammatic and sentential representations (i.e., Cheng, 2002; Larkin & Simon, 1987) or compared the performance of reasoners using spatial representations with the performance of reasoners using sentential representations (i.e., Bauer & Johnson-Laird, 1993; Gick & Holyoak, 1983; Pedone, Hummel, & Holyoak, 2001).

* Present address: School of Psychology, Cardiff University.
E-mail address: gattism@cardiff.ac.uk (M. Gattis).

0364-0213/$ – see front matter © 2004 Cognitive Science Society, Inc. All rights reserved.
doi:10.1016/j.cogsci.2004.02.001

Both types of studies have tried to pinpoint how the reasoning process is influenced by the representational format of a reasoning problem.

This paper addresses the question of how spatial representations influence reasoning by focusing on a different problem: how spatial representations acquire conceptual meaning. This question is pertinent to spatial reasoning because some spatial representations improve reasoning, and some do not, and one reason for such failures is that some representations do not convey their intended meanings. A good theory of spatial reasoning ought to explain why some spatial representations work and others do not, and also be able to predict which types of representations are in the future likely to benefit reasoning. In addition, while many have argued that spatial representations benefit reasoning by making relations explicit (Bauer & Johnson-Laird, 1993; Gattis & Holyoak, 1996; Larkin & Simon, 1987), relatively little has been said about how they do so. The experiments presented here test one hypothesis about how spatial representations make conceptual relations explicit: that reasoners interpret spatial representations by creating an analogical mapping between the conceptual and spatial relations based on similarities in relational structure. First, however, I would like to briefly review some older ideas about how spatial representations acquire meaning (see Gattis, 2001, for a more detailed review).

1. How spatial representations acquire meaning

For the most part, symbols acquire meaning through *learned pairings*. Decoding written language, for instance, involves learning the pairings between script and sound, which makes learning to read a long and laborious process (Goswami & Bryant, 1990; Treiman, 2000). Pictures, diagrams, and other spatial representations may acquire meaning through learned pairings (how else would anyone understand Kandinsky?!), but unlike written language, spatial representations such as diagrams, graphs, and gestures frequently acquire meaning through other avenues as well. Though traditionally the domain of semiotics, this question has also been studied by psychologists and linguists, and work in all three areas has suggested at least two ways in which spatial representations acquire meaning other than through learned pairings.

Iconicity influences the interpretation of spatial representations whenever a spatial representation maintains some physical resemblance to the thing it represents (Fromkin & Rodman, 1998; McNeill, 1992; Peirce, 1903/1960; Tversky, 1995). An iconic representation is isomorphic rather than identical to its object. Partial resemblance allows flexibility in the representation, and it comes as no surprise that iconicity is a powerful and ubiquitous representational tool. The iconicity of many spatial representations is transparent, particularly maps and diagrams of physical systems such as gears, planes, and levers (for examples, see Funt, 1980; Hegarty, 1992; Larkin & Simon, 1987; Mayer & Gallini, 1990; Narayan, Suwa, & Motoda, 1995; Schwartz, 1995). Iconicity suffers, however, from two important limitations. Peirce claimed that icons cannot convey any new information, since the only thing an icon can do is represent an object that already exists. Additionally, because iconicity relies on physical resemblance, it is unable to represent abstract qualities or concepts, such as “goodness” or “value.”

Because of the limitations of iconicity in spatial representations, abstract qualities and concepts are frequently represented by association with some object. *Associations* influence spatial representations whenever a correspondence is established between a depicted object and a concept or quality that is not depicted but is associated by past experience with what is depicted (Fromkin & Rodman, 1998; Tversky, 1995, 2001; Werner & Kaplan, 1963/1984). Associations then do not rely on resemblance of appearance but on experience. Associations between physical aspects of the world and conceptual aspects of experience are frequently reflected in language, such as the association between "more" and "up" reflected in metaphorical expressions like "My income rose last year" (Lakoff & Johnson, 1980). This association is seen in graphs as well, for example, in the use of height on the vertical axis to represent "more." Children and adults in a variety of cultures exhibit a tendency to rely on an association between "up" and "more" in both the construction and the interpretation of spatial representations (Tversky, Kugelmass, & Winter, 1991; see also Handel, DeSoto, & London, 1968). It is easy to see how such associations can be derived from experience: for solid objects on a solid surface, for instance, greater quantity indeed entails greater vertical extent. Associations allow spatial representations to move beyond depiction to represent abstractions. Warm or cool temperatures can be represented by objects associated with those temperatures, such as fire or ice; fast or slow rates can be represented by race cars or snails; and happy or sad moods can be represented by spatial metaphors like "feeling high" and "feeling low." Unlike the specific experience of learned pairings underlying symbols, associations rely on general experience and as a result are generalizable to new representations.

Together iconicity and associations provide powerful and flexible bases for mapping concepts to visuospatial elements when interpreting spatial representations. There are several reasons to think that these two semiotic principles are inadequate, however, for explaining how spatial representations acquire meaning. We are able to create and interpret an amazingly wide variety of novel spatial representations, ranging from calendars and graphs to the box-and-arrow diagrams of many scientists, and many other less conventionalized and novel forms. Our ability to communicate with spatial representations relies on consensus between a wide variety of users without specific instruction concerning the pairings or mappings between the components of the representation and the concepts being represented, and despite the varying past experiences that users bring to the task. What is surprising about the way that humans reason with spatial representations is that some spatial representations acquire consistent meanings, even though those consistencies that are not readily explained by similarities based on appearance or past experience. For instance, when asked to interpret Cartesian line graphs, both adults and young children with no graphing experience infer that steeper lines represent faster rates of change and less steep lines represent slower rates of change (the visual equivalent of the graphing rule that the rate of change equals $\Delta y/\Delta x$) (Gattis, 2002; Gattis & Holyoak, 1996). Associations do not provide a satisfactory explanation of this finding because our experience in the physical world includes traveling downhill, when steeper is indeed faster, and traveling uphill, when steeper is painfully slower. In addition, a spatial representation may potentially contain multiple mappings, and when multiple mappings conflict, some mappings reliably take precedence over others. Gattis and Holyoak (1996) asked adults to reason with graphs that contrasted the iconic mapping of "up" on a vertical line and "up" in the atmosphere against the mapping between steeper slope and faster rate of change. The mapping between slope and rate exerted a stronger influence on reasoning performance than the mapping based on verticality. Thus some additional principle appears to guide which mapping is used.

2. Mapping relational structure

A third basis for how spatial representations may acquire meaning is that the assignment of conceptual meaning to spatial representations may be based on similarities of relational structure. *Relational structure* may influence the assignment of meaning to spatial representations when assignments, or mappings, are based on correspondences between objects, relations between them, and relations between relations. The relational structure hypothesis predicts that in the interpretation of spatial representations, conceptual elements are mapped to spatial elements, conceptual relations are mapped to spatial relations, and higher-order conceptual relations are mapped to higher-order spatial relations.

One reason to think that relational structure may influence reasoning with spatial representations is that similarities of relational structure influence many other cognitive processes, including both reasoning about concepts (such as metaphors or analogical reasoning) and comparisons between perceptual stimuli (Gentner, 1977, 1983, 1988; Gentner & Rattermann, 1991; Goldstone, Medin, & Gentner, 1991; Markman & Gentner, 1993). Conceptual analogies such as "An electric battery is like a reservoir" or "The atom is like the solar system" function by activating our knowledge not just about similar objects, but also about similar relations between objects (Gentner, 1983; Holyoak & Thagard, 1989). Relational similarity also influences the second step in analogical reasoning, mapping aspects of the two concepts together. An analogy or metaphor may be based on a mapping between the attributes or features of objects (sometimes called an attribute mapping or attribute match), or on a mapping between the relations between objects (sometimes called a relational mapping or relational match) (Markman & Gentner, 1993). Young children tend to create mappings between objects based on shared attributes, such as "A tiger is like a zebra"; as they get older, children increasingly create mappings based on shared relational structure, such as "Tree bark is like skin" (Gentner, 1988). Similarly, when asked to rate the aptness of metaphors, adults give higher ratings to relational metaphors such as, "A tire is a shoe" and lower ratings to appearance-based metaphors such as, "The sun is an orange" (Gentner, 1988).

Relational similarity also influences judgments about perceptual stimuli. In a pictorial analogy task, Gentner (1977) showed preschoolers, first-graders, and adults pictures of mountains and trees, and asked them questions like "If this tree had a knee, where would it be?" and "If this mountain had a chin, where would it be?" Children in both age groups identified body parts in a way that maintained consistent spatial relations between the parts, so that for instance the vertical order eyes, nose, mouth, chin was more likely than the vertical order nose, mouth, eyes, chin when mapping facial features to a mountain, despite the fact that appearance matches (eyelike features on a mountain, for instance) could lead to a vertical order that was inconsistent with the vertical order on the human face. Experiments with more complex perceptual stimuli have shown that both attribute matches and relational matches influence mapping between perceptual stimuli, and similarity judgments of perceptual stimuli, but that the contribution of attributes and relations to mapping or to perceived similarity depends upon the relative strength of each (Goldstone et al., 1991; Markman & Gentner, 1993).

One of the few studies to examine how similarities of relational structure may influence the mapping *between* conceptual and perceptual stimuli investigated how children first begin to reason with graphs. Gattis (2002) asked young children with no prior graphing experience to reason about differences in quantity or differences in rate using graph-like diagrams that resembled Cartesian line graphs. Children's answers to questions like "Does this line stand for more or less?" and "Does this line stand for faster or slower?" revealed two consistent mappings of concepts to space. Quantity judgments were consistent with the height of a line and rate judgments were consistent with the slope of a line. Quantity and height are both first-order relations (relations between elements) and rate and slope are both second-order relations (relations between relations), suggesting that young children are not only sensitive to similarities of relational structure within concepts and within perceptual stimuli, but between concepts and percepts as well.

The studies reported here test the hypothesis that similarities of relational structure influence reasoning with diagrams, so that between concepts and percepts elements are mapped to elements, and relations are mapped to relations. These experiments focus on very simple relational structures—objects and relations between objects—to explore the role of relational structure in spatial reasoning. Objects and relations were chosen for two reasons. First, as described above, the distinction between objects and relations has received a great deal of attention in studies of conceptual and studies of spatial analogies, and the results of those studies suggest that the mapping of objects to objects and relations to relations plays an important role in comparisons between concepts and in comparisons between perceptual stimuli, but do not tell us about mappings between concepts and perceptual stimuli. Second, the results of Gattis (2002) provided evidence that similarities of relational structure influence the mapping of different levels of relations between concepts and perceptual stimuli, but did not address whether similarities of relational structure influence the mapping of objects and relations between concepts and perceptual stimuli.

In order to test the relational structure hypothesis in a representational system that would be as novel as possible for the participants (and thus eliminate the chance that participants had already learned certain pairings), an experimental procedure was developed involving diagrams of gestures. Three experiments used diagrams of gestures to investigate whether adults' interpretations of novel spatial representations would be influenced by similarities of relational structure. If spatial reasoning is based on similarities of relational structure, judgment patterns ought to reflect mapping of conceptual elements to spatial elements and conceptual relations to spatial relations.

3. Experiment 1: relational structure in active declarative statements

Experiment 1 examined whether relational structure influences the interpretation of diagrams representing active declarative statements. Active declarative statements were chosen because they are basic linguistic structures (Fromkin & Rodman, 1998), and because they allow two different relations to be contrasted in a simple way. Adults were asked to interpret diagrams of hand gestures in a simple and quick test. In the first phase, two drawings, each of a character extending one hand, were paired with sentences assigning an animal to each hand, such as "This hand means MOUSE" and "This hand means BEAR" (see Fig. 1A). In the second phase, two new drawings were paired with two active declarative statements, both of which involved the meaning assigned to the right hand. The two drawings were of the same character touching

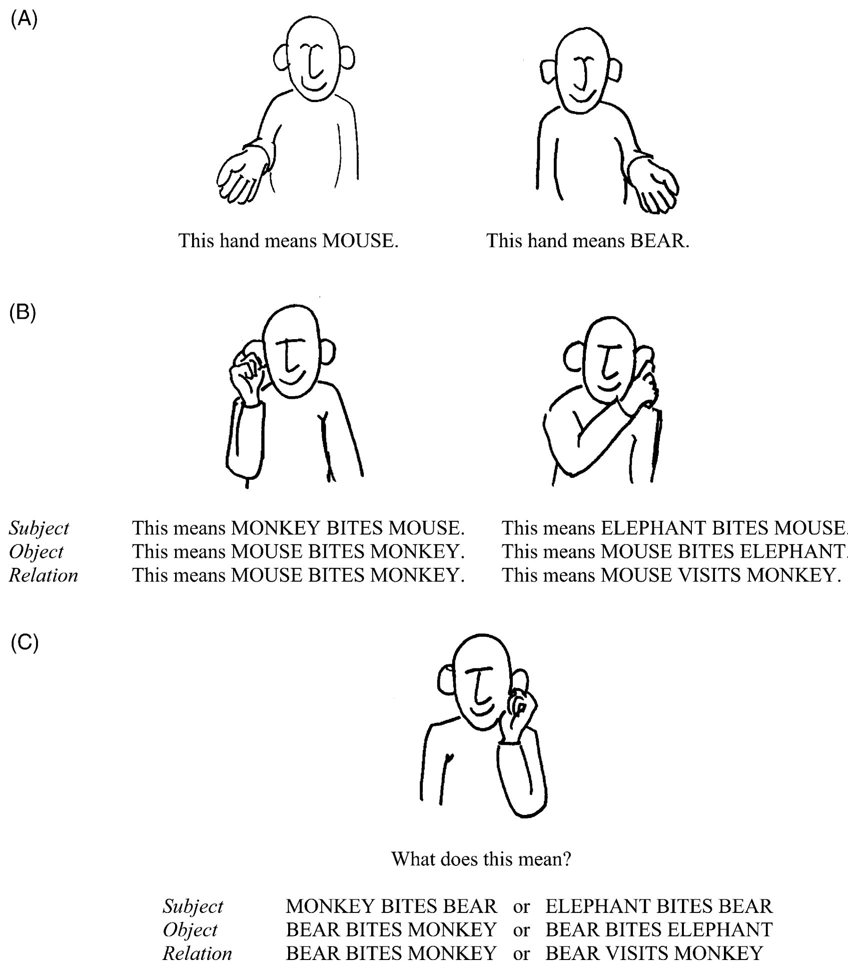

Fig. 1. Examples of the materials given to participants. In the first phase of Experiments 1 and 2 participants were given diagrams paired with an initial mapping (A). Participants in Experiments 1 and 3 were given diagrams paired with active declarative statements in which either the subject, the object, or the relation varied (B). The final step for all experiments was a forced-choice judgment about the meaning of two new diagrams (C).

his right ear with his right hand and touching his left ear with his right hand, as seen in Fig. 1B. The statements were about animal characters performing some action toward another animal character, and either the subject, the relation, or the object varied between the two statements. For example, when the subject varied, the statements might have been “Monkey bites Mouse” and “Elephant bites Mouse.” The statements introduced two new animals, in this case Monkey and Elephant, both of whom perform the same action: they bite Mouse. When the object varied, the statements might have been “Mouse bites Monkey” and “Mouse bites Elephant.” As in the subject-varying condition, in the object-varying condition, the new statements introduced two new animals, here again Monkey and Elephant, but now the new animals are the objects of the action rather than the subjects of the action. When the relation varied, the statements might have been "Mouse bites Monkey," and "Mouse visits Monkey." In the relation-varying condition, the two new statements introduced just one new animal, here Monkey, who was the recipient of two different actions, biting and visiting. In all three of the experimental conditions (subject varying, object varying, and relation varying), the assignment of meaning to the diagram was intentionally ambiguous. As can be seen in Fig. 1B, assigning one statement to each sign left open whether it was the object touched by the hand (right ear, left ear) or the relation of the arm to body (ipsilateral, contralateral) that carried that particular meaning.

The third phase introduced gestures complementary to those seen earlier, but made with the left hand rather than the right (touching the left ear with the left hand and touching the right ear with the left hand, illustrated in Fig. 1C). These gestures were used to probe how people assigned meaning to the diagrams—whether to the object touched by the hand or to the relation of the hand to the body—by asking participants to judge which of two new statements was represented by each new sign. The statements involved previously introduced animals and actions in novel combinations that varied according to condition but corresponded to the statements in the previous phase. In the subject-varying condition described above, the probes would have been "Monkey bites Bear" and "Elephant bites Bear." In the object-varying condition described above, the probes would have been "Bear bites Monkey" and "Bear bites Elephant." In the corresponding relation-varying condition, the probes would have been "Bear bites Monkey" and "Bear visits Monkey."

The experimental design allowed a specific hypothesis about spatial reasoning to be tested. If reasoners are influenced by similarities of relational structure when interpreting diagrams, it should be the case that meaning is mapped to the diagrams differently depending on whether the varying portion of the statement was an object (one of the animal characters, regardless of whether that character was the subject or object of the statement) or a relation (the actions biting and visiting). For example, if a reasoner in the varying-object condition chose "Bear bites Elephant" as the meaning of the left-hand-touching-left-ear sign, and "Bear bites Monkey" as the meaning of the "left-hand-touching-right-ear sign, by comparing that assignment with the diagram-statement pairings in Phase 2, we can conclude that Elephant was assigned to the left ear, and Monkey was assigned to the right ear (see Fig. 2A). More generally we can conclude that the varying portion of the statements, the animal characters, were assigned to the two ears, reflecting a mapping of objects to objects. This type of mapping will be called an object-mapping. By contrast, if a reasoner in the same condition chose "Bear bites Monkey" as the meaning of the left-hand-touching-left-ear sign, and "Bear bites Elephant" as the meaning of the "left-hand-touching-right-ear sign", again by comparing those choices with the previously assigned statements, we can conclude that the Monkey was assigned to the straight or ipsilateral arm, and Elephant to the crossed or contralateral arm (see Fig. 2B). More generally we can conclude that the varying portion of the statements were assigned to the relation of the arm to the body. This type of mapping will be called a relation-mapping.

The relational structure hypothesis thus predicts that people will make object-mappings in conditions in which objects are varying, and relation-mappings in the condition in which relations are varying. Alternatively, it may be that similarities of relational structure do not influence the interpretation of diagrams, or do not outweigh other influences. One possibility is that people use individual strategies to try to assign meaning to diagrams. If individual strategies are numerous and varied, and no single alternate strategy dominates the mapping

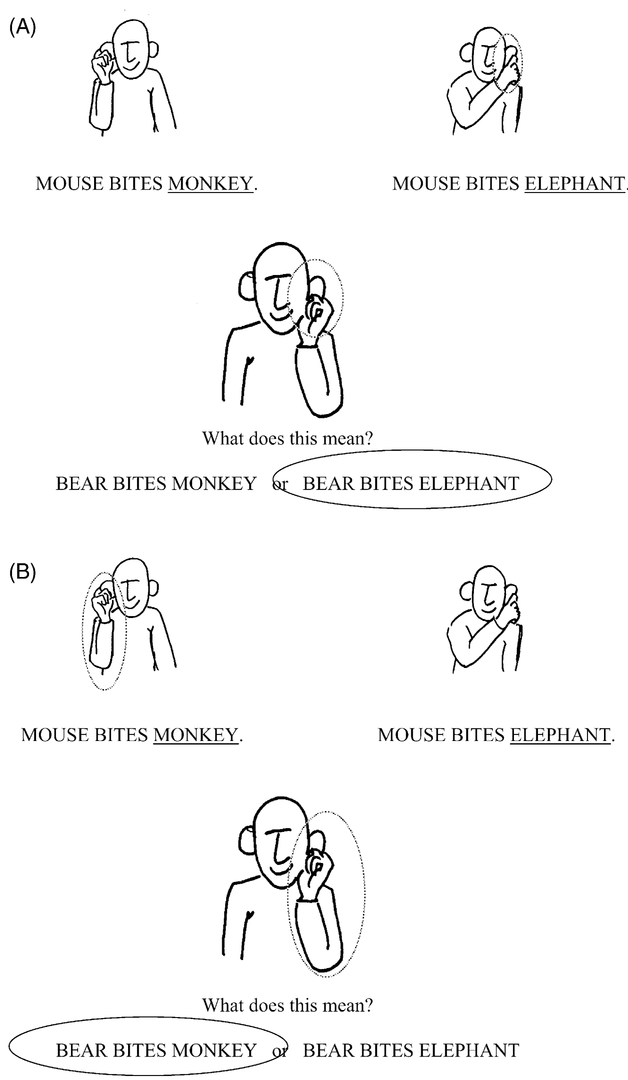

Fig. 2. Example of an object-mapping (A) and a relation-mapping (B).

process, those differences can be expected to cancel one another out when the results are pooled across participants, and we should not expect the results to differ from chance. If a single strategy other than relational similarity dominates, we can expect to see some other pattern of mapping emerge when the results are pooled across participants.

3.1. Method

3.1.1. Participants

One hundred and ninety-two students from the University of Munich and the Technical University of Chemnitz participated in Experiment 1 during psychology classes. Participation was voluntary. Approximately one-third of the students were randomly assigned to each of the three experimental conditions.

Any participant who gave the same answer to both probe questions was discarded from the analyses, with the justification that these two questions were on the same page, and served as a consistency check. This rule was applied in all experiments reported here.

3.1.2. Procedure and design

Participants were given a booklet of three sheets of A4 paper stapled together. The experiment was conducted in German, and the following description provides an English translation of the materials. Each page contained two diagrams, each accompanied by a simple declarative statement. Each diagram and the accompanying statement occupied approximately half a page, and the materials were organized vertically so that the first diagram occupied the top half of the first page, the second diagram occupied the bottom half of the first page, and so on. Thus it was easy to compare the two study examples, and the two test items, but participants were not encouraged to turn back to the examples at test. The instructions were, "Please read the following carefully. At the end you will be asked questions about it." No specific time constraints were given but participants were encouraged to perform the task quickly, and all participants completed the experiment within a minute or two.

On the first page of the booklet were the two drawings shown in Fig. 1A, and below each was a sentence of the form "This hand means (Animal)." A random order of four animals—Monkey, Elephant, Mouse, and Bear—was created, and three additional orders were created by rotating the list, so that for some participants Elephant was the first animal introduced, for some Mouse, and for some Bear. These four orders were counterbalanced between subjects. At the bottom of the first page was an instruction, "Please turn to the next page."

On the second page of the booklet were the two drawings shown in Fig. 1B. The order of these two drawings was counterbalanced between subjects. Above each drawing was a sentence, and the two sentences differed in either the subject (subject-varying), the object (object-varying), or the action performed (relation-varying). This between-subjects variable was the main experimental manipulation. For the subject-varying condition, the two sentences were of the form, "This means 'Animal3 R-action Animal1' " and "This means 'Animal4 R-action Animal1.' " The relation (R-action) was either "visits" or "bites" and was counterbalanced between subjects. For the object-varying condition, the two sentences were of the form, "This means 'Animal1 R-action Animal3' " and "This means 'Animal1 R-action Animal4,' " and again the relation was either "visits" or "bites" and was counterbalanced across subjects. For the relation-varying condition, the two sentences were of the form, "This means 'Animal1 R-action1 Animal3' " and "This means 'Animal1 R-action2 Animal3.' " As with the subject-varying condition, the relations were "visits" and "bites," and the order of these two relations was counterbalanced between subjects. At the bottom of the page was the instruction, "Please turn to the next page."

On the third page of the booklet were two drawings of the same character touching his left ear with his left hand, and touching his right ear with his left hand (the first of these is illustrated in Fig. 1C). The order of the two drawings was counterbalanced. Above each drawing was the question, "What does this mean?", and below each drawing were two sentences, with the instruction, "Circle the answer that fits best." For the subject-varying condition, the two sentences were of the form, "This means 'Animal3 R-action Animal2' *OR* This means 'Animal4

Table 1
Proportions of object-mappings and relation-mappings for each of the experimental conditions in Experiment 1

Condition	Proportion of participants choosing a mapping		Total number of participants
	Object-mapping	Relation-mapping	
Subject-varying	64.3	35.7	56
Object-varying	72.6	27.4	62
Relation-varying	38.7	61.3	62

R-action Animal2.' " For the object-varying condition, the two sentences were of the form, "This means 'Animal2 R-action Animal3' *OR* This means 'Animal2 R-action Animal4.' " For the relation-varying condition, the two sentences were of the form, "This means 'Animal2 R-action1 Animal3' *OR* This means 'Animal2 R-action2 Animal3.' " For each condition, the order of the two sentences was counterbalanced between subjects. Note that the two sentences are identical to those introduced in the second phase, with the single change that Animal2, which has already been assigned to the left hand, is substituted for Animal1, to correspond to the left-handed actions in the drawings.

The experimental design thus involved one experimental variable with three levels, plus five counterbalancing variables with four, two, two, two and two levels each. This meant that there were 192 conditions in all, with one participant in each condition.

3.2. Results and discussion

Asking participants to judge the meaning of two signs with the left hand allowed for a consistency measure, in that the two signs logically ought to have two different meanings. This requirement seemed especially appropriate given that both test items were on the same page, allowing an easy comparison between the two. Thus if participants shifted the basis for mapping between the first and second test items, it was easy for them to change their first answer (and some did). On this basis, 12 participants who chose the same meaning for both diagrams were discarded from the analyses. The answers given by the remaining 180 participants were then coded by whether the unassigned meaning was mapped to the ears (object-mapping) or to the ipsilateral and contralateral relations of the arm to the body (relation-mapping), as described above. The frequency of object-mappings and relation-mappings were then compared. The mapping pattern varied significantly between conditions $\chi^2(2, N = 180) = 15.82$, $p < .01$. As can be seen in Table 1, participants in the subject-varying condition and the object-varying condition were more likely to choose meanings that reflected object-mappings, whereas participants in the relation-varying condition were more likely to choose meanings that reflected relation-mappings. Mean proportions for each of the possible mappings (verbal object to spatial object, verbal object to spatial relation, verbal relation to spatial relation, and verbal relation to spatial object) were as follows: 68.5 object–object, 31.5 object–relation, 61.3 relation–relation, and 38.7 relation–object.

The results of Experiment 1 were consistent with the predictions of the relational structure hypothesis. When the subject or object of the active declarative statement was unassigned, participants tended to map the varying objects to physical objects—the right and left ears. When the relation of the active declarative statement was unassigned, participants tended to map the varying relations to physical relations—the ipsilateral and contralateral relation of the arm to the body.

4. Experiment 2: relational structure in conjunctive and disjunctive statements

Experiment 2 investigated whether the similarities of relational structures would also influence the interpretation of diagrams representing a very different type of relational structure: conjunctions and disjunctions. Conjunctions and disjunctions play an important role in deductive reasoning (Johnson-Laird & Byrne, 1991), in probability judgments (Osherson, 1995), and in causal reasoning (Eels, 1991; Östermeier & Hesse, 2000), so it seems important to understand how these relations might be understood diagrammatically. Conjunctions and disjunctions are also good examples of abstract relations that are difficult to represent via iconicity or associations. Perhaps as a result, logicians have created symbols for these important relations, "and" being represented as $\cap$, and "or" being represented as $\cup$—a classic example of symbols that acquire meanings through learned pairings rather than some other more direct route. Understanding how conjunctions and disjunctions might be represented diagrammatically thus has implications for theories of reasoning, as well as for education and graphic design.

Experiment 2 used the same diagrams of gestures and the three phase procedure used in Experiment 1, with the difference that diagrams were paired with simple conjunctions and disjunctions of animal characters, such as "Monkey and Mouse," and "Monkey or Mouse." As in Experiment 1, different types of relational structure were contrasted by varying different aspects of the statement for different participants. In the second and third phases of the experiment, either the first animal (O1), the second animal (O2), or the relation between them (R) varied. The hands always signified the first two animals introduced, but neither the relations nor subsequent animals were specifically assigned to any aspect of the diagram.

Again the hypothesis was that if the similarity of relational structures influences diagram interpretation, reasoners would map the unassigned and varying portion of the statement with a structurally similar aspect of the accompanying sign, so that objects are mapped to objects and relations are mapped to relations. Varying animals, whether in the first position or the second position, were expected to be mapped to the ears. Varying relations, in this case "and" and "or," were expected to be mapped to the ispilateral and contralateral relations of the arm to the rest of the body. These two mapping patterns were again predicted to lead to opposite judgments in the final phase.

4.1. Method

4.1.1. Participants

One hundred and ninety-three students from the University of Munich and the Technical University of Chemnitz participated in Experiment 2 during psychology classes. Participation was voluntary. Approximately one-third of the students were randomly assigned to each of the three conditions.

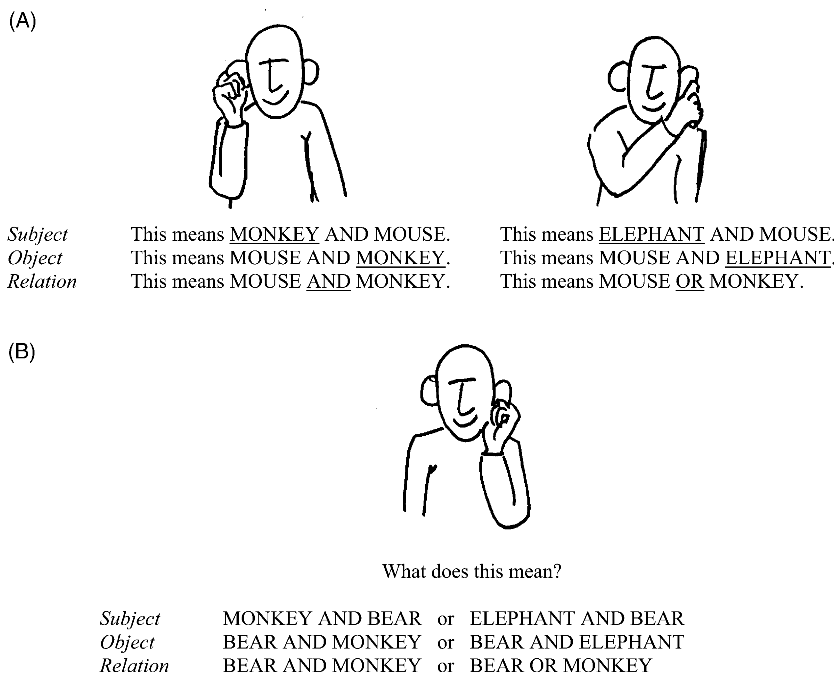

Fig. 3. Example of the drawings and conjunctive and disjunctive statements (A) and judgment tasks (B) given to participants in each of the three conditions of Experiment 2.

Because of the fact that each time we conducted the experiment a small number of students did not return the experimental materials, we always handed out slightly more booklets than the number required for our counterbalancing conditions. Sometimes this meant that we had more participants than our design required, as in this experiment, when we ended up with one extra participant in the object-varying condition.

4.1.2. Procedure and design

The procedure and materials were nearly identical to those of Experiment 1. There were three experimental conditions: the first animal in a statement varied (O1), the second animal in a statement varied (O2), or the relation between them varied (R). The relations were “and” and “or,” and the order of these two relations was counterbalanced between subjects. Examples of each of these conditions can be seen in Fig. 3A. Statements in the probe phase corresponded to the conditions described for Experiment 1, as shown in Fig. 3B.

4.2. Results and discussion

Fifteen participants who indicated that both signs had the same meaning were excluded from the analyses. The answers given by the remaining 178 participants were then coded by whether the unassigned meaning was mapped to the ears (object-mapping) or to the ipsilateral and contralateral relations of the arm to the body (relation-mapping), as described

Table 2
Proportions of object-mappings and relation-mappings for each of the experimental conditions in Experiment 2

Condition	Proportion of participants choosing a mapping		Total number of participants
	Object-mapping	Relation-mapping	
O1-varying	69.0	31.0	58
O2-varying	65.6	34.4	61
Relation-varying	32.2	67.8	59

above. The frequency of object-mappings and relation-mappings were then compared. The mapping pattern varied significantly between conditions, $\chi^2(2, N = 178) = 19.73$, $p < .01$. As can be seen in Table 2, participants in the O1 condition and the O2 condition were more likely to make object-mappings, whereas participants in the R condition were more likely to make relation-mappings. Mean proportions for each of the possible mappings were as follows: 67.3 object–object, 32.7 object–relation, 67.8 relation–relation, and 32.2 relation–object.

The results of Experiment 2 were consistent with the predictions of the relational structure hypothesis and provided evidence that relational structure can be identified in new contexts with some consistency. The results were similar to the results of Experiment 1 in that when asked to interpret novel diagrams of gestures, people tended to map elements to elements and relations to relations. The results of Experiment 2 extended the findings of Experiment 1 in an important way, by showing that mapping of relational structures occurs in a variety of syntactical contexts, including conjunctive and disjunctive statements, thus ruling out the possibility that the use of verbs alone bias people toward a relational mapping.

5. Experiment 3: relational structure without initial mappings

Together the results of Experiments 1 and 2 indicate that people interpret diagrams of gestures by establishing correspondences between relational structures contained in the diagrams and relational structures contained in the meanings assigned to them. Much like analogical reasoning, this sensitivity to similarities of relational structure in spatial reasoning leads people to map objects to objects and relations to relations, and to extend this mapping to subsequent interpretations of novel diagrams. In both experiments, however, participants were given an object-to-object mapping in the first step of the experiment, when they were told that each hand stood for a particular animal (see Fig. 1). On this basis it could be argued the results of Experiments 1 and 2 demonstrated that people reasoning with diagrams are sensitive to clues or instruction about initial mappings, and are able to extend on those initial mappings by making subsequent judgments that are consistent with the initial mappings provided, but that diagrams may acquire meaning via relational structure only when some instruction about an initial mapping is provided. A stronger test of the robustness of relational similarity would involve asking people to judge the meanings of diagrams without providing them with an initial object-to-object mapping. Experiment 3 did just that.

Experiment 3 was identical to Experiment 1, with the difference that the first phase of Experiment 1 (the initial mapping) was not included. This served as a manipulation check to test whether the structurally-sensitive mappings observed in Experiments 1 and 2 were due to the initial mapping. Experiment 3 had just two short steps. First people were given diagrams of gestures paired with active declarative statements. The diagrams and statements were the same as those used in the second phase of Experiment 1, examples of which can be seen in Fig. 1B. As in Experiment 1, different types of relational structure were contrasted by varying different aspects of the statement for different participants, and either the subject, the object, or the relation of the statements varied. In the second phase of the experiment people were given two new diagrams and were asked to make a forced choice about the meaning of those two diagrams.

The hypothesis was the same as in the previous two experiments: if the similarity of relational structures influences spatial reasoning, independent of the presence or absence of instruction establishing relational correspondences, reasoners ought to map the unassigned and varying portion of the statement with a structurally similar aspect of the accompanying sign, so that objects are mapped to objects and relations are mapped to relations. These two mapping patterns were predicted to lead to opposite judgments about the meaning of new diagrams in the final phase.

5.1. Method

5.1.1. Participants

One hundred and ninety-four students from the University of Sheffield participated in Experiment 3 during psychology classes. Participation was voluntary. Approximately one-third of the students were randomly assigned to each of the three conditions. As in Experiment 2 we had slightly more than the required 192 participants, this time two extra participants in the relation-varying condition.

5.1.2. Procedure and design

The design, procedure, and materials were nearly identical to those of Experiment 1, except that the experiment was conducted in English, and the experimental materials consisted of only two pages, which would have been the second and third phases of Experiment 1. There were three experimental conditions: the subject of each statement varied (subject-varying), the object of each statement varied (object-varying), or the active declarative relation in each statement varied (relation-varying). Examples of the first step for each of the conditions can be seen in Fig. 1B. Examples of the second step for each of the conditions can be seen in Fig. 1C.

5.2. Results and discussion

Four participants who indicated that both signs had the same meaning were excluded from the analyses. The answers given by the remaining participants were then coded by whether the unassigned meaning was mapped to the ears (object-mapping) or to the ipsilateral and contralateral relations of the arm to the body (relation-mapping), as described above. The frequency of object- and relation-mappings were then compared. The mapping pattern varied significantly

Table 3
Proportions of object-mappings and relation-mappings for each of the experimental conditions in Experiment 3

Condition	Proportion of participants choosing a mapping		Total number of participants
	Object-mapping	Relation-mapping	
Subject-varying	61.9	38.1	63
Object-varying	60.9	39.1	64
Relation-varying	36.5	63.5	63

between conditions, $\chi^2(2, N = 190) = 10.51$, $p < .01$. As can be seen in Table 3, participants in the subject-varying condition and the object-varying condition were more likely to make object-mappings, whereas participants in the relation-varying condition were more likely to make relation-mappings. Mean proportions for each of the possible mappings were as follows: 61.4 object–object, 38.6 object–relation, 63.5 relation–relation, and 36.5 relation–object.

6. General discussion

This paper concerns how spatial representations acquire meaning and whether the mapping of relational structures might play a role in that process. Understanding how spatial representations acquire meaning is pertinent to cognitive scientists and philosophers interested in theories of representation and to designers and educators interested in communicating accurately and efficiently. Three experiments tested the hypothesis that diagrams may acquire meaning through a basic inferential reasoning process, that of aligning similar relational structures contained in the diagram and in the concepts being reasoned about, and using the resulting correspondences to infer missing information.

In all three experiments, adults were shown novel diagrams of gestures paired with simple statements and either the subject, the relation, or the object of the statements varied. Experiment 1 tested whether the mapping of relational structure influences the interpretation of diagrams paired with active declarative statements. Experiment 2 tested whether the mapping of relational structure occurs in a variety of syntactical contexts by pairing the diagrams with conjunctive and disjunctive statements. Experiment 3 tested whether the mapping of relational structures is dependent on an explicit initial mapping by using the same procedure as Experiment 1, but without the initial mapping provided in Experiments 1 and 2. All three experiments involved a final step in which participants judged the meaning of new gestures, and those judgments were used to assess how people assign meaning to spatial representations.

Despite the simple and humorous style of the diagrams and the statements, the procedure did resemble many situations in which people reason with spatial representations, in that participants were given a general meaning for each diagram, but were not told what particular aspect of a diagram carried a particular meaning. This ambiguous assignment of meaning to the diagram was used to test the relational structure hypothesis, which predicted that when the subject or object of the statements varied, people would assign the unmapped meaning to an object in the diagram and when the relation of the statements varied, people would assign the unmapped meaning to a spatial relation in the diagram. The results of all three experiments are consistent with the hypothesis that similarity of relational structures influences spatial reasoning: the majority of people chose an interpretation that indicated a mapping of objects to objects and of relations to relations, in two different syntactical contexts and with and without an explicit initial mapping.

6.1. Relational structure and the interpretation of spatial representations

The results reported here are consistent with the hypothesis that relational structure influences the interpretation of spatial representations, but the limitations of the current evidence must be acknowledged. Across the three experiments, in conditions in which objects varied, 66% of people chose object-mappings and in conditions in which relations varied, 65% of people chose relation-mappings. One question that thus arises is the generalizability of the current evidence, and more specifically how the remaining 1/3 of people solved the problem of meaning assignment. Because the task was designed to be novel and to provide minimal guidance about the representational system being introduced, one realistic possibility is that some participants did not feel the task provided a basis for an informed response, and therefore simply guessed.

A second and more interesting possibility is that mapping relational structure is just one of several constraints or sources of similarity influencing diagram interpretation. Indeed, as discussed in the introduction, at least two other constraints in diagram interpretation are well-known: iconicity and associations. Related studies using the same diagrammatic materials paired with spatial prepositions such as near/far and above/below indicate that iconicity influences diagram interpretation on this task at least some of the time (Gattis, 2003). In those studies, approximately the same proportion of people chose object-mappings in object-varying conditions, but the number of people choosing relation-mappings in relation-varying conditions dropped to chance. Gattis (2003) argued that some participants may have been influenced by an iconic mapping of the prepositions and others by a relational mapping of the prepositions leading to two different types of responses in the relation-varying conditions. In addition, performance on spatial reasoning tasks in general may be an interaction between similarity and individual strategies (Gattis & Dupeyrat, 2000, 2003). Future studies might compare different sources of similarity in problems, and consider participants' justifications of their interpretations to investigate individual strategies and competing sources of similarity further.

The question of generalizability also raises the issue of whether relational structure influences reasoning with a variety of spatial representations or only the diagrams described here. The experiments reported in this paper tested the relational structure hypothesis with just one set of diagrams, and to avoid influences of prior knowledge of learned conventions, the diagrams were novel, and were designed specifically for these experiments. Several other converging pieces of evidence, however, suggest that similarities of relational structure may influence reasoning with a variety of spatial representations. The use of elements to represent elements and relations to represent relations may be found in many conventionalized and familiar diagrams, such as flow charts, organizational charts and Venn diagrams. In addition, the claims of the relational structure hypothesis are consistent with several other reports of diagrammatic and spatial reasoning (e.g., Bauer & Johnson-Laird, 1993; Gattis, 2002; Gattis

& Holyoak, 1996; Issing, 1990; Novick & Hurley, 2001; Novick, Hurley, & Francis, 1999; Tversky, 2001). In a study comparing different assignments of variables to axes in Cartesian line graphs, Gattis and Holyoak (1996) found that people reasoned more accurately about rate when the queried variable was assigned to the vertical axis—independent of whether that variable was the independent variable or the dependent variable. This variable assignment creates a low-level mapping between "more" and "up," and a higher-level mapping between steeper function lines and faster rates of change. The primacy of higher-level mappings over low-level mappings parallels a similar result in analogical reasoning (Gentner, 1983), suggesting that relational structure exerts similar influences on analogical reasoning and reasoning with graphs.

A more recent study of the biases children demonstrate in graphical reasoning provides further support for the hypothesis that the mapping of relational structure influences spatial reasoning. When Gattis (2002) asked young children to reason about quantity or rate using graph-like diagrams, children tended to map similar levels of relational structures together, mapping quantity to the height of a line and rate to the slope of a line, even though they had no prior experience with graphs. The results of that study thus extend the claims of this paper to more realistic type of spatial representation, and to more complex relations.

Novick and Hurley (2001) also reported that when math and science students are asked to interpret or choose an appropriate representation for a particular set of information, their choices and their justification of their choices indicate that they are sensitive to the structure of a problem, and are aware of how different problem structures are best represented diagrammatically. The current results demonstrate that such sensitivity to structure is not only found with familiar spatial representations, but with novel spatial representations as well, suggesting that structural sensitivity in spatial reasoning may be a product of basic cognitive processes, not just learned convention. Future research might investigate the relational structure hypothesis using diagrams and spatial representations other than Cartesian graphs and the gesture diagrams presented here. For instance, many novel visual representations with varying degrees of expressive power can be found in handbooks of visual representation (e.g., Tufte, 1983, 1990, 1997), in math and science textbooks, and in the work of cognitive scientists, designers and educators who are interested in designing new representational systems (e.g., Cheng, 1999, 2002).

6.2. Analogies and diagrams

From the early likening of sound and waves to the more recent comparison of armies and rays, many analogies intertwine spatial and conceptual components so tightly that it seems difficult to unravel how they first came together. Perhaps this melding is one reason why many investigations of analogy have involved comparison of problems that may be presented verbally or visually without addressing how a theory of analogies and diagrams are specifically related (e.g., Beveridge & Parkins, 1987; Gick & Holyoak, 1983; Pedone et al., 2001). Other studies of analogical reasoning have used diagrams or drawings to present a story, as in Markman and Gentner's (1993) picture-stories of a man giving food to a woman and a woman giving food to a squirrel. One reason why diagrams can be powerful source analogs in analogical reasoning may be that our perceptual system is organized around objects and the relations between them, leading in turn to a representation that emphasizes relational structures. Whereas these and other studies of similarity and analogical reasoning have demonstrated that relational structure influences mapping and comparison processes of concepts *and* of perceptual stimuli, however, the work presented here demonstrates that relational structure influences mapping of concepts *to* perceptual stimuli. The results reported in this paper thus indicate that relational structure plays an important role in how diagrams are invested with meaning.

The results reported here also suggest several directions for future research concerning the extent to which the same cognitive processes underlie both spatial and analogical reasoning. Studies of analogical reasoning have traditionally distinguished several steps in the process, including encoding, retrieval, and mapping (e.g., Holyoak, Novick, & Melz, 1994; Holyoak & Thagard, 1989, 1995; Ross, 1989). Such studies have shown that some failures in analogical reasoning are due to failures to encode, others to failures to retrieve, and others to failures to map. Those findings suggest an alternate explanation to that discussed above for why the performances of approximately one-third of the people participating in these experiments did not conform to the predictions of the relational structure hypothesis: some people may have failed to encode the distinction between objects and relations at all, and as a result been unable to map according to that distinction. In order to test the hypothesis of interest here, the task was designed with a minimum of examples and guidance about how to perform the task: participants saw just two examples before test, were encouraged to perform the task quickly, and were not encouraged to turn back to the examples at test (though it was not possible to prevent participants from doing so, given that they were tested in large numbers). Future studies of diagrammatic reasoning might consider designs in which the number of examples and the opportunity to refer back to the examples at test are manipulated (for instance, by presenting the items on computer rather than paper) in order to investigate whether failures are occurring at encoding, retrieval, or mapping.

6.3. Language and space

The question of how spatial and conceptual information are linked via structure mapping may have implications not only for understanding the use of spatial representations in reasoning, but also for aspects of language such as gesture (e.g., Goldin-Meadow, McNeill, & Singleton, 1996; Kita, Danziger, & Stolz, 2001; McNeill, 1992), naming (e.g., Bloom & Markson, 1998), verb learning (Fisher, 1996), and more generally for understanding how spatial structures influence and are influenced by the structure of languages (e.g., Bloom, Peterson, Nadel, & Garrett, 1996). Relational language seems to play an important role in the development of analogical reasoning, helping children to notice similarities of relational structure and use them in analogy-making (Gentner & Rattermann, 1991; Loewenstein & Gentner, 2001). Research with signed languages also indicates that in signing space, objects or actors are assigned to a spatial locus (Emmorey, 1996, 2001). One interesting question for future research is to what extent the representational rules underlying diagrammatic representation and signed languages are similar or dissimilar.

Studies of linguistic structure also suggest one alternative to the relational structure hypothesis, namely that people are mapping roles and movement rather than structure per se when interpreting diagrams (e.g., Lakoff & Johnson, 1980). For instance, the ipsilateral and contralateral gestures could be interpreted as movements rather than bodily relations. The tendency to pair the actions "bites" and "visits" with the ipsilateral and contralateral relations could be seen to emphasize the movement toward a location or animal, rather than a bodily relation. This explanation is related to the idea that diagrams acquire meaning via associations (discussed in the introduction of this paper), because it postulates that people associate movement of the arms with movement to a location, or the movement of an action such as "bite" or "visit." If people perceive and map the movement path, however, we would also expect that the movement marks the grammatical roles of subject and object, as in signed languages (Emmorey, 1996, 2001). For example, when movement is used to represent an action in American Sign Language, such as "The dog bites the cat," the direction of the movement marks the grammatical roles of subject and object: the subject is the starting location, and the object is the end location. Were adults simply mapping locations and actions to the signs shown here by mapping location and action to a movement path, we would expect to find stronger mapping patterns for those situations in which the object of the statement was mapped to the ears compared to those in which the subject of the statement was mapped to the ears. In Experiments 1 and 3, however, subjects and objects of active declarative statements were mapped with equal frequency to the ears, indicating that perceived movement did not play an important role in the interpretations of diagrams presented here.

6.4. Iconicity, associations, and relational structure

This paper has addressed the question of how spatial representations acquire conceptual meaning, and has discussed three constraints or sources of similarity influencing that process. The appeal of two of these constraints, iconicity and associations, has been discussed by psychologists and semioticians for many years, and these two factors indeed play a role in diagram interpretation. The results of the experiments reported here indicate that a third constraint, relational structure, also plays a strong role in diagram interpretation and more generally suggest why spatial representations can be such powerful tools for reasoning. The relations between these three constraints are an unexplored and promising area for future research.

As discussed earlier, one possibility to explore is whether individuals differ in their reliance on or sensitivity to these different constraints. A related but independent question is how iconicity, associations, and relational structure contribute to particular spatial representations and particular reasoning tasks. More specifically, each constraint influences different aspects of the concepts involved: iconicity emphasizes physical resemblance, associations emphasize attributes, and relational structure emphasizes relations between objects and the relations between relations. One simple prediction that follows is that the representations most likely to benefit reasoning are those that convey meaning via a constraint that corresponds to the reasoning task by emphasizing or activating task-appropriate aspects of the concept.

A third aspect of the relations between iconicity, associations, and relational structure for future work to explore is how these constraints unfold developmentally. The results of several studies indicate that children are sensitive to all three constraints by the early school years. For instance, Werner and Kaplan (1963/1984) and Tversky et al. (1991) discussed the influence of iconicity and associations on young children's production and interpretation of diagrams and graphs. Gattis (2002) also reported that 6- and 7-year-olds were sensitive to similarities of relational structure in interpreting line graphs. Studies with much younger children are required to understand better the emergence of these constraints and the relations between them.

6.5. Conclusion

This paper proposes that the similarity of relational structures influences not just mapping within concepts or within percepts, but also between concepts and percepts. The findings reported here (and also related findings in Gattis, 2002) support the relational structure hypothesis, suggesting that relational structure may be an important process for the construction and interpretation of spatial representations. By introducing a new paradigm for studying the mapping of meaning to space, the experiments reported here may provide an interesting task for studying relational structure in reasoning and language. The diagrams used in these studies involve simple but important relational structures, and offer a method for directly testing how conceptual structures are mapped to spatial structures. Further research might use this paradigm to address the relational structure underlying other logical and linguistic structures as well.

Acknowledgments

This work was conducted at the Max Planck Institute for Psychological Research and the University of Sheffield, and was funded by the Max Planck Society. I thank Marc Bühner, Rob Goldstone, and Michael Waldmann for useful discussions in the initial stages of developing these experiments, and all the scientists at the Max Planck Institute who participated in the pilot studies and made helpful suggestions. I thank Christophe Lardschneider for making the drawings used in these experiments, Denise Parks for assistance in creating materials and conducting experiments at the University of Munich, and Martin Baumann, Marc Bühner, and Andreas Keinath for conducting experiments at the University of Technology, Chemnitz and the University of Sheffield. Preliminary reports of Experiments 1 and 2 were presented at the Workshop on Advances in Analogy Research at New Bulgarian University, Sofia, Bulgaria, July 1998, and at Mind: The Annual Conference of the Cognitive Science Society of Ireland, Dublin, Ireland, August 1998.

References

Bauer, M. I., & Johnson-Laird, P. N. (1993). How diagrams can improve reasoning. *Psychological Science, 4*, 372–378.

Beveridge, M., & Parkins, E. (1987). Visual representation in analogical problem solving. *Memory* & *Cognition, 15*, 230–237.

Bloom, P., & Markson, L. (1998). Intention and analogy in children's naming of pictorial representations. *Psychological Science, 9*, 200–204.

Bloom, P., Peterson, M., Nadel, L., & Garrett, M. (1996). *Language and space*. Cambridge, MA: MIT Press.

Cheng, P. C. H. (1999). Interactive law encoding diagrams for learning and instruction. *Learning and Instruction, 9*, 309–325.

Cheng, P. C. H. (2002). Electrifying diagrams for learning: Principles for complex representational systems. *Cognitive Science*, *26*, 685–736.

Eels, E. (1991). *Probabilistic causality*. Cambridge: Cambridge University Press.

Emmorey, K. (1996). The confluence of space and language in signed languages. In P. Bloom, M. Peterson, L. Nadel, & M. Garrett (Eds.), *Language and space* (pp. 171–209). Cambridge, MA: MIT Press.

Emmorey, K. (2001). Space on hand: The exploitation of signing space to illustrate abstract thought. In M. Gattis (Ed.), *Spatial schemas in abstract thought* (pp. 147–174). Cambridge, MA: MIT Press.

Fisher, C. (1996). Structural limits on verb mapping: The role of analogy in children's interpretation of sentences. *Cognitive Psychology*, *31*, 41–81.

Fromkin, V., & Rodman, R. (1998). *An introduction to language* (6th ed.). Fort Worth, TX: Harcourt Brace.

Funt, B. V. (1980). Problem-solving with diagrammatic representations. *Artificial Intelligence*, *13*, 201–230.

Gattis, M. (2001). Mapping conceptual and spatial schemas. In M. Gattis (Ed.), *Spatial schemas in abstract thought* (pp. 223–245). Cambridge, MA: MIT Press.

Gattis, M. (2002). Structure mapping in spatial reasoning. *Cognitive Development*, *17*, 1157–1183.

Gattis, M. (2003). How similarity shapes diagrams. *Lecture Notes in Artificial Intelligence*, *2685*, 249–262.

Gattis, M., & Dupeyrat, C. (2000). Spatial strategies in reasoning. In W. Schaeken, A. Vandierendonck, & G. de Vooght (Eds.), *Deductive reasoning and strategies* (pp. 153–175). Hillsdale, NJ: Erlbaum.

Gattis, M., & Dupeyrat, C. (2003). *Spatial schemas in everyday reasoning*. Manuscript in preparation.

Gattis, M., & Holyoak, K. J. (1996). Mapping conceptual to spatial relations in visual reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 231–239.

Gentner, D. (1977). Children's performance on a spatial analogies task. *Child Development*, *48*, 1034–1039.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.

Gentner, D. (1988). Metaphor as structure mapping: The relational shift. *Child Development*, *59*, 47–59.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. Gelman & J. P. Byrnes (Eds.), *Perspectives on thought and language: Interrelations in development*. London: Cambridge University Press.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, *15*, 1–38.

Goldin-Meadow, S., McNeill, D., & Singleton, J. (1996). Silence is liberating: Removing the handcuffs on grammatical expression in the manual modality. *Psychological Review*, *103*, 34–55.

Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the nonindependence of features in similarity judgments. *Cognitive Psychology*, *23*, 222–262.

Goswami, U., & Bryant, P. (1990). *Phonological skills and learning to read*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Handel, S., DeSoto, C. B., & London, M. (1968). Reasoning and spatial representations. *Journal of Verbal Learning and Verbal Behavior*, *7*, 351–357.

Hegarty, M. (1992). Mental animation: Inferring motion from static displays of mechanical systems. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *18*, 1084–1102.

Holyoak, K. J., Novick, L. R., & Melz, E. R. (1994). Component processes in analogical transfer: Mapping, pattern completion, and adaptation. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory, Vol. 2: Analogical connections* (pp. 113–180). Norwood, NJ: Ablex.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, *13*, 295–355.

Holyoak, K. J., & Thagard, P. (1995). *Mental leaps: Analogy in creative thought*. Cambridge, MA: The MIT Press.

Issing, L. J. (1990). Learning from pictorial analogies. *European Journal of Psychology of Education*, *4*, 489–499.

Johnson-Laird, P. N., & Byrne, R. M. J. (1991). *Deduction*. Essays in Cognitive Psychology Series. London: Lawrence Erlbaum Associates.

Kita, S., Danziger, E., & Stolz, C. (2001). Cultural specificity of spatial schemas, as manifested in spontaneous gestures. In M. Gattis (Ed.), *Spatial schemas in abstract thought* (pp. 115–146). Cambridge, MA: MIT Press.

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.

Loewenstein, J., & Gentner, D. (2001). Spatial mapping in preschoolers: Close comparisons facilitate far mappings. *Journal of Cognition and Development*, *2*, 189–219.

Markman, A. B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology*, *25*, 431–467.

Mayer, R. E., & Gallini, J. K. (1990). When is an illustration worth ten thousand words? *Journal of Educational Psychology*, *82*, 715–726.

McNeill, D. (1992). *Hand and mind: What gestures reveal about thought*. Chicago: University of Chicago Press.

Narayan, N. H., Suwa, M., & Motoda, H. (1995). Behavior hypothesis from schematic diagrams. In J. Glasgow, N. H. Narayan, & B. Chandrasekaran (Eds.), *Diagrammatic reasoning: Cognitive and computational perspectives* (pp. 501–534). Menlo Park, CA: AAAI Press/The MIT Press.

Novick, L. R., & Hurley, S. M. (2001). To matrix, network, or hierarchy: That is the question. *Cognitive Psychology*, *42*, 158–216.

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

Osherson, D. (1995). Probability judgment. In E. E. Smith & D. N. Osherson (Eds.), *An invitation to cognitive science, Vol. 3: Thinking* (pp. 35–75). Cambridge, MA: The MIT Press.

Östermeier, U., & Hesse, F. W. (2000). Verbal and visual causal arguments. *Cognition*, *75*, 65–104.

Pedone, R., Hummel, J. E., & Holyoak, K. J. (2001). The use of diagrams in analogical problem solving. *Memory* & *Cognition*, *29*, 214–221.

Peirce, C. S. (1960). In C. Hartshorne & P. Weiss (Eds.), *Collected papers, Vol. II: Elements of logic*. Cambridge, MA: The Belknap Press of Harvard University Press. (Original work published 1903).

Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 456–468.

Schwartz, D. L. (1995). Reasoning about the referent of a pictures versus reasoning about the picture as the referent: An effect of visual realism. *Memory* & *Cognition*, *23*, 709–722.

Treiman, R. (2000). The foundations of literacy. *Current Directions in Psychological Science*, *9*, 89–92.

Tufte, E. R. (1983). *The visual display of quantitative information*. Chesire, CT: Graphics Press.

Tufte, E. R. (1990). *Envisioning information*. Chesire, CT: Graphics Press.

Tufte, E. R. (1997). *Visual explanations*. Chesire, CT: Graphics Press.

Tversky, B. (1995). Cognitive origins of graphic conventions. In F. T. Marchese (Ed.), *Understanding images* (pp. 29–53). New York: Springer-Verlag.

Tversky, B. (2001). Spatial schemas in depictions. In M. Gattis (Ed.), *Spatial schemas in abstract thought* (pp. 79–112). Cambridge, MA: MIT Press.

Tversky, B., Kugelmass, S., & Winter, A. (1991). Cross-cultural and developmental trends in graphic productions. *Cognitive Psychology*, *23*, 515–557.

Werner, H., & Kaplan, B. (1984). *Symbol formation: An organismic-developmental approach to the psychology of language*. Hillsdale, NJ: Lawrence Erlbaum Associates. (Original work published 1963).