Simulating the Effects of Relational Language in the Development of Spatial Mapping Abilities

Tom Mostek[1], Jeff Loewenstein[2], Ken Forbus[1], Dedre Gentner[2]

Northwestern University

1890 Maple Avenue

Evanston, IL, 60201, USA

{t-mostek, loewenstein, forbus, gentner} @nwu.edu

Abstract

Young children's performance on certain mapping tasks can be improved by introducing relational language (Gentner, 1998). We show that children's performance on a spatial mapping task can be modeled using the Structure-Mapping Engine (SME) to simulate the comparisons involved. To model the effects of relational language in our simulations, we vary the quantity and nature of the spatial relations and object descriptions represented. The results reproduce the trends observed in the developmental studies of Loewenstein & Gentner (1998; in preparation). The results of these simulations are consistent with the claim that gains in relational representation are a major contributor to the development of spatial mapping ability. We further suggest that relational language can promote relational representation.

Introduction

Spatial reasoning is one of the core abilities in human cognition. An important test of spatial reasoning is the *mapping task* (DeLoache, 1987, 1995; Huttenlocher, Newcombe, & Sandberg, 1994; Uttal, Schreiber, & DeLoache, 1995; Uttal, Gregg, Tan, Chamberlin, & Sines, submitted). In a mapping task, the goal is to find a correspondence between two different spatial situations. In DeLoache's classic task, a child is shown two rooms, similar in layout and furniture (though not necessarily in size). A toy is hidden in one room and the child must look for another toy in the corresponding place in the other room (e.g., DeLoache, 1995). It has been proposed (Gentner & Rattermann, 1991) that the same process of structural alignment that is used in analogy and similarity may play a role in spatial mapping tasks. That is, spatial mapping tasks can be viewed as a kind of analogy in which the spatial relationships of the situations involved provide the base and target descriptions for the structural alignment, and the correspondences computed in structural alignment provide the basis for inferring the correct answer.

This paper provides evidence for the role of structural alignment in spatial mapping tasks. We show how the pattern of developmental results found by Loewenstein and Gentner (1998, in preparation) can be modeled using SME (the Structure-Mapping Engine (Falkenhainer, Forbus, & Gentner, 1989; Forbus, Ferguson, & Gentner 1994)) a simulation of Gentner's (1983) structure-mapping theory. We start by describing Loewenstein and Gentner's spatial mapping task. Next we describe how we used SME to model the results.

Spatial Mapping Tasks

Mapping and symbolic reference is ubiquitous in adult daily life, but it develops only gradually in children. Studies by Blades and Cooke (1994), DeLoache (1995), Uttal (Uttal, Schreiber, & DeLoache, 1995), and others have shown that preschool children have great difficulty with the seemingly simple task of finding an object in the 'same place' as an object in an almost identical model, even though they can easily retrieve the original hidden object. Gentner and her colleagues have suggested that one contribution to the great gains children make in their performance on spatial mapping tasks is relational knowledge, and further, that acquiring relational language promotes this relational knowledge (Gentner & Loewenstein, in preparation; Gentner & Rattermann, 1991; Gentner, Rattermann, Kotovsky, & Markman, 1995; Kotovsky & Gentner, 1996; Loewenstein & Gentner, 1998).

Experiment 1. The first study (Loewenstein & Gentner, 1998) used the setup in the left of Figure 1, with neutral appearances for the cards. Three cards are placed *on*, *in*, and *under* the Hiding box. The instructions given in the baseline condition avoided language that used spatial relationships. During the orientation trial, the experimenter said "I'm putting the winner right here" while placing the card in its location at the Hiding box. The instructions in the language condition used spatial relationships during the orientation task to describe where the cards were being placed: The experimenter said, "I'm putting the winner [in/on/under] the box." For both conditions, no language was used in the finding task: The Experimenter gestured generally towards the Finding box, saying "Can you find the

---

[1] Computer Science Department
[2] Psychology Department

Neutral Objects

Hiding Box   Finding Box

Relational Match

Cross-Mapped Objects

Hiding Box   Finding Box

**Object Matches Compete
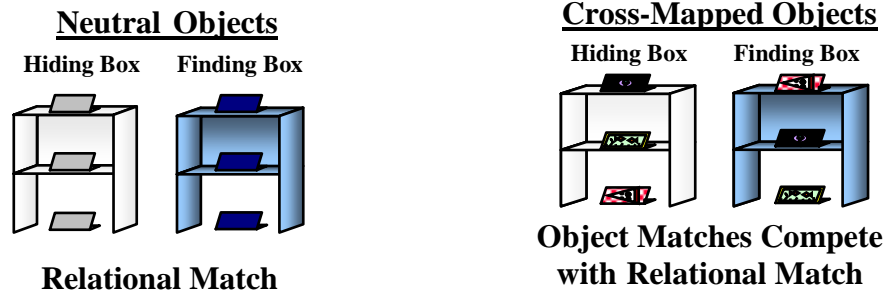with Relational Match**

Figure 1 – Experiment Setup

winner here, in the very same place?" In the baseline condition, 44-month-olds found the sticker only 42% of the time, not significantly above chance performance of 33%. However, in the language condition, they performed far better, finding the winner 70% of the time (see Figure 2, left). The 49-month-olds performed fairly well in both the baseline (63%) and language (73%) conditions. It appears that hearing spatial relational terms led the younger children to form stronger representations of the spatial relational structure.

Experiment 2. Gentner and Loewenstein (in preparation) use a similar task as in Experiment 1, but the three cards associated with a box were all distinctive and unique, as shown in Figure 1. While there is an exact object match for each card, the cards that match in appearance have different spatial relationships with the box. For example, the card that is ON the Hiding box matches the card that is IN the Finding box. This is an example of a *cross-mapping task* (Gentner & Toupin, 1986; Gentner & Rattermann, 1991), in which object similarity is pitted against relational similarity. Such tasks are useful in testing for the availability and salience of the child's relational knowledge. The results bore out previous findings that cross-mapping tasks are difficult: 49-month-olds were at chance at finding the winner in both conditions. Even 62-month-olds were correct only 53% of the time in the baseline condition. However, in the language condition, when given the spatial relation during the hiding task, their performance improves to 73% correct (see Figure 2, right).

These results suggest the following conjectures:
- Age-related improvements are largely due to improved understanding of spatial relationships.
- Relational language highlights spatial relationships, supporting children's relational mapping abilities.

This pattern of results is consistent with other findings on the role of relational language in domain learning (e.g., Gentner & Rattermann, 1991; Gentner, Rattermann, Markman & Kotovsky, 1995; Kotovsky & Gentner, 1996), and lends evidence to the position that relational language fosters the development of relational thought (Gentner, 1998).

Modeling Spatial Mapping as Visual Comparisons

The spatial mapping task above involves encoding descriptions of the hiding box and the finding box and comparing these descriptions to predict, based on the location of the winner in the hiding box, where the winner will be in the finding box. Since we are modeling the comparison process via structure-mapping, we first briefly review structure-mapping theory and SME.

Review of Structure-Mapping

According to structure-mapping theory, the process of structural alignment takes as input two structured representations (*base* and *target*) and produces as output a set of
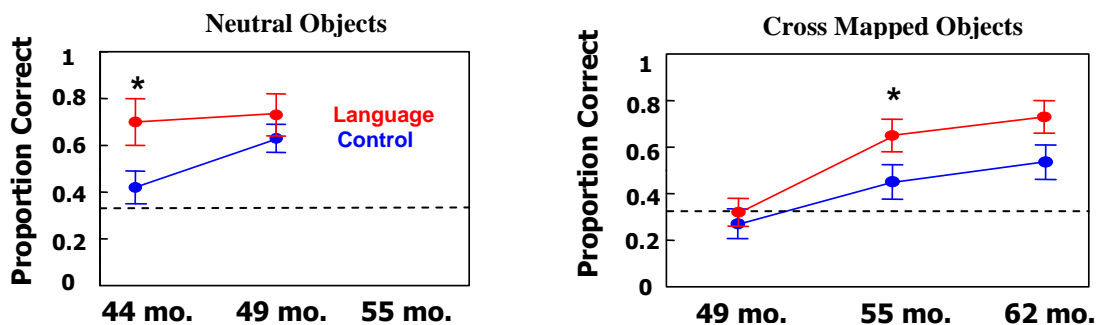

Figure 2 – Experiment Results

mappings. Each mapping consists of a set of *correspondences* that aligns items in the base with items in the target and a set of *candidate inferences*, which are surmises about the target made on the basis of the base representation plus the correspondences. The constraints on the correspondences include *structural consistency,* i.e., that each item in the base maps to at most one item in the target and vice-versa (the *1:1 constraint*), and that if a correspondence between two statements is included in a mapping, then so must correspondences between its arguments (the *parallel connectivity* constraint). Which mapping is chosen is governed by the *systematicity* constraint: Preference is given to mappings that match systems of relations in the base and target. Each of these constraints is motivated by the role analogy plays in cognitive processing. The 1:1 and parallel connectivity constraints ensure that the candidate inferences are well-defined. The systematicity constraint reflects a (tacit) preference for inferential power in analogical arguments.

The Structure-Mapping Engine (SME) (Falkenhainer *et al* 1989; Forbus *et al* 1994) is a cognitive simulation of analogical matching. Given base and target descriptions, SME finds globally consistent interpretations via a local-to-global match process. SME begins by proposing correspondences, called *match hypotheses,* in parallel between statements in the base and target. Then, SME filters out structurally inconsistent match hypotheses. Mutually consistent collections of match hypotheses are gathered into global mappings using a greedy merge algorithm. An evaluation procedure based on the systematicity principle is used to compute the *structural evaluation* for each match hypothesis and mapping. These numerical estimates are used both to guide the merge process and as one component in the evaluation of an analogy. It is important to note that SME can produce multiple mappings for a given pair of base and target descriptions, corresponding to different ways in which they might be aligned.

Using SME to Model the Effects of Relational Knowledge in Spatial Mapping

Our focus is on demonstrating that, if SME is used to model structural alignment, that changes in available relational knowledge can explain the pattern of results. Consequently, we do not model the encoding processes themselves, only their results. Our goal is to show that the assumed outcomes of encoding in the explanations for the experiments do in fact lead to the pattern of results found, given that SME is used to model the structural alignment involved.

We model the role of structural alignment in these tasks using SME by the following assumptions:
- The base is our construal of what the child might have encoded about the hiding box.

- The target is our construal of what the child might have encoded about the finding box.
- The child's response will be based on the mapping they use, specifically, the candidate inference for which card has the sticker behind it (and hence is the "winner").
- The use of relational language leads to increased relational content in the child's representations
- The likelihood that a child uses a particular mapping is a function of its structural evaluation.

This last assumption requires further explanation. Why not just take SME's top-rated mapping as the child's mapping, as was done in Markman & Gentner (1993) or Gentner, Rattermann, Markman, and Kotovsky (1995)? We believe that variability in children's performance is caused by variations in how they encode the situations. There are many sources of encoding variability, including domain knowledge and context. We model the potential variability in these processes by representing what the final result of such processes might be, and taking alternate mappings generated by SME as representative of different possible outcomes of an interleaved encoding/mapping process.

For each experiment, we generated a set of descriptions intended to be representative of the encodings children used at different ages and different conditions. We then ran SME on these descriptions to compute the mappings. The prediction made by a mapping is generated from the candidate inferences for a mapping: The location of the card in the Hiding box is represented by the statement (BEHIND *<card>* STICKER) in the base, with no such statement appearing in the target. The lack of corresponding statement in the target means that there will be a candidate inference computed from this statement. Consequently, whatever corresponds to the winning card in the hiding box within a mapping will be the predicted winning card in the Finding box. We randomly chose a location for the winner card, and this location had no effect on the results.

As noted above, SME can produce multiple mappings for a given base and target. We used these multiple mappings to estimate the relative likelihood of a particular card being suggested as the winner for each base and target as follows. To eliminate size effects, the score for each mapping is divided by an ideal score (i.e., the score obtained by mapping the target to itself). The relative likelihood of choosing a card is calculated by the sum of the scores of all the mappings that predict it. The relative likelihoods for the cards are scaled so that they sum to 1.

Since the specific values of these numbers depend on the particular choice of representations and processing parameters, they must be interpreted with care. Specifically, we do not consider the particular numerical values produced as the probability that a specific outcome will occur. For robustness, we only consider meaningful the ordinal relationships between these numbers. That is, if one number calculated is larger/smaller than another, then the corres-

```
(defdescription exp1-target        (at box card-3)                    (open-object card-2)
  :documentation                                                      (folding-card card-2)
  "neutral object finding task"    ;; unique relations                (horizontally-oriented card-2)
  :entities                        (on box card-1)                    (green-colored card-2)
   (box card-1 card-2 card-3)      (in--cont-open box card-2)
  :expressions                     (under box card-3)                 (open-object card-3)
 ((open-object box)                                                   (folding-card card-3)
  (rectangle-volume box)           ;; cards                           (horizontally-oriented card-3)
  (vertically-oriented box)        (open-object card-1)               (green-colored card-3)))
  (touches box ground)             (folding-card card-1)
  (at box card-1)                  (horizontally-oriented card-1)
  (at box card-2)                  (green-colored card-1)
```

Figure 3 – Sample Description

ponding frequency in the human data should be larger/smaller.

Simulation Experiment 1

To model the effects of relational knowledge in Experiment 1 we used three different base descriptions. All base descriptions used the spatial relation AT (e.g., (AT BOX CARD-1)), descriptions of the box and the three cards, and the location of the winner (i.e., (BEHIND CARD-1 STICKER)). The rest of the base content varied as follows:
- Level 1: No additional information
- Level 2: Includes the spatial relationship between the box and the winner (e.g., (IN CARD-2 BOX)).
- Level 3: Included all spatial relations involving the cards and the boxes.

In other words, with each increase in level is an increase in the number of first-order spatial relations.

The target description included descriptions of the box, the cards, and all of the spatial relationships involving them, but did not mention the location of the sticker. (It is sufficient to vary the base representations without varying the target representations because the failure to include a proposition in either base or target will prevent a correspondence from being formed.)

We then ran SME on each of the three base descriptions and the target. For each pair, SME generates multiple legitimate mappings (e.g. the first mapping has the top card corresponding to the top card, and the second mapping has the top card corresponding to the middle card). We used the procedure outlined above to calculate the accuracy that each level of description predicts.

Results

As expected, increasing the number of unique spatial relations resulted in stronger mappings between the cards 'in the same place' in each representation. This in turn resulted in a greater calculated probability of the correct card being selected as 'the winner'. With no specific spatial relations represented (Level 1), the model gave each possible mapping the same score. (yielding 33% selection probability for the correct relational choice). With one spatial relation representing the location of the winner (Level 2), the model gave the correct choice the greatest score (40%). When all three first-order spatial relations were represented (Level 3), the model's selection probability for the correct choice increased (to 48%). Thus, the model predicts (quite reasonably) that children who did not represent any spatial relations would perform at chance (33%), and that more fully representing the spatial relations in the two scenes would increase performance.

Discussion

In Lowenstein and Gentner's study, the 44 month olds performed at chance in the baseline condition and at 70% in the language condition. The simulation results suggest that in the baseline condition, the children were not encoding any of the spatial relations that uniquely identify the card. The language condition explicitly draws attention to one of the identifying spatial relations. The simulation results suggest that encoding this information is sufficient to raise

	Level 1			Level 2			Level 3		
	Score	% from best	Selection probability	Score	% from best	Selection probability	Score	% from best	Selection probability
'the winner'	0.224	0	33.3%	0.241	N/A	40%	0.266	N/A	48%
Other card 1	0.224	0	33.3%	0.232	3.7	30%	0.252	5.7	26%
Other card 2	0.224	0	33.3%	0.232	3.7	30%	0.252	5.7	26%

Table 1 – Results from Simulation Experiment 1
Note – max score in these cases is 0.302

the performance as seen. In the second simulation experiment we consider the more difficult cross-mapped model task and investigate the effect of encoding higher-order relations.

Simulation Experiment 2

Recall that Experiment 2 pitted relational similarity against object similarity. The rich object matches between the individual cards suggest a different set of correspondences than the parallel spatial relationships that define the correct responses. We suggested that the effects of age and relational language in increasing accuracy both stemmed from greater encoding of relations (thereby preventing strong attribute overlap from overwhelming relational mappings). Our purpose in this simulation was to see if varying the amount of relational structure does indeed lead to these effects.

The method used is similar to that used in modeling Experiment 1 – the target representation is fully specified, and the base representation is varied with respect to its relational specificity. The cards are described at the same level of detail in both target and base descriptions. In contrast to Experiment 1 (for which the object descriptions were all identical and could be represented with only four attributes), the rich cross-mapped objects used in Experiment 2 required nine attributes and two relations. They were represented as uniquely colored and as having a different picture on the front of each object. The base descriptions include different levels of information, as follows:

3. Basic relations – The relations ON, IN and UNDER are included. (the Level 3 case in Experiment 1)
4. Extra binary relations – In addition to ON, IN and UNDER, several binary relations are added, specifically (ABOVE CARD-1 CARD-2) and (ABOVE CARD-2 CARD-3)
5. Extra ternary relation – In addition to ON, IN and UNDER, a relation that ties all three cards together is added: (IN-A-COLUMN CARD-1 CARD-2 CARD-3)
6. One higher-order relation – In addition to ON, IN and UNDER, some of the inferential structure linking relations is added: (IMPLIES (AND (ON CARD-1 BOX) (IN CARD-2 BOX)) (ABOVE CARD-1 CARD-2))
7. Two higher-order relations – Like #4, but with two higher-order relations: (IMPLIES (AND (IN CARD-2

BOX) (UNDER CARD-3 BOX)) (ABOVE CARD-2 CARD-3))

We then ran SME using each of these bases with the same target, as we did in the first model, and generated relative predictions using the same method.

Results

The results range from a below chance selection probability of 27% (always preferring the 'same card' attribute mapping over the 'same place' card) to a strong performance of 60%, as shown in Table 2. The model's selection probability for the correct choice increases with increased relational knowledge. This is consistent with the possibility that children's improved performance between 49 and 62 months of age, and between the baseline and language condition, is due to better encoding of the spatial relations relevant for the task.

In order for the model to consistently make the correct relational matches in the face of richly represented cross-mapped objects, higher-order relations needed to be included in the representations. Even with all the basic first-order relations represented (Level 3) -- the most successful model in Simulation Experiment 1 -- the model still chose the object choice (47% selection probability for the object match). Adding a ternary relation was also not sufficient to enable the model to make the correct relational choice. One higher-order relation was just enough to induce a shift to the relational choice as the mapping with the best score. However, the scores of the relational mapping and the object mapping were extremely close (relational choice: 50%; object match: 47%). Representing two higher-order relations yields a clearly dominant selection of the correct relational choice (60%). Although the precise threshold at which relational overlap will overcome attribute overlap depends on specific modeling parameters (such as the amount of each type of information available), what is clear here is that the addition of higher-order relations contributes to improved performance. Thus, the improved performance of children in Experiment 2 may be due to learning of higher-order relations.

General Discussion

| | Level 3 | | | Level 4 | | | Level 5 | | | Level 6 | | | Level 7 | | |
	Score	% from best	Prob.	Score	% from best	Prob.	Score	% from best	Prob.	Score	% from best	Prob.	Score	% from best	Prob.
Winner	4.50	16.5	27%	4.55	16.4	27%	4.52	16.3	27%	7.77	N/A	50%	8.36	N/A	60%
Same	5.39	N/A	47%	5.44	N/A	47%	5.40	N/A	47%	7.67	1.3	47%	7.67	8.2	35%
Other	4.48	16.7	26%	4.53	16.6	26%	4.49	16.7	26%	6.32	18.8	3%	6.85	18.0	5%

Table 2 – Results from Simulation Experiment 2
Note – max score in these cases is 9.67

Relational language helped preschool children perform mapping tasks with neutral and cross-mapped objects. However, because the effects came about at different ages, it was possible that language was providing a different kind of support for 44- and 62-month-olds. The simulation experiments provide evidence for this suggestion by showing a link between level of relational understanding and mapping task performance. The simulation experiments also provide support for the claim that spatial mapping tasks involve structural alignment. Simulation experiment 1, with no competing object mappings, showed that first order relations are sufficient for success on the relational mapping task. Simulation Experiment 2, with rich cross-mapped objects, suggested that higher-order relations are needed to perform relational mappings under these more challenging circumstances. Further experimentation providing higher-order relations could provide evidence for this hypothesis.

During spatial mapping tasks, as in many everyday tasks, there is a vast array of perceptual information available to the child. Our evidence suggests that knowing what to encode may be in part a learned skill. Learning relational language may contribute to this ability.

Acknowledgements

References

Blades, M., & Cooke, Z. (1994). Young children's ability to understand a model as a spatial representation. *Journal of Genetic Psychology, 155,* 201-218.

DeLoache, J. S. (1987). Rapid change in the symbolic functioning of very young children. *Science*, *238*, 1556-1557.

DeLoache, J. S. (1989). Young children's understanding of the correspondence between a scale model and a larger space. *Cognitive Development, 4,* 121-139.

DeLoache, J. S. (1995). Early symbol understanding and use. In D. Medin (Ed.) *The psychology of learning and motivation* (Vol. 32). New York: Academic Press.

Falkenhainer, B., Forbus, K., & Gentner, D. (1989). The Structure Mapping Engine: Algorithm and Examples. *Artificial Intelligence, 41.* (pp. 1-63).

Forbus, K. D., Ferguson, R. W., & Gentner, D. (1994). Incremental structure mapping. In A. Ram & K. Eiselt (Eds.*), Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* (pp. 313-318). Atlanta, GA:Lawrence Erlbaum Associates.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science. Vol 7(2)*, (pp. 155-170).

Gentner, D. (1998). Relational language and relational thought. Plenary paper presented at the Meeting of the Piaget Society, Chicago, IL.

Gentner, D., & Loewenstein, J. (in preparation). Relational language contributes to the development of spatial mapping.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes (Eds.), *Perspectives on Language & Thought: Interrelations in Development* (pp. 255-277). Cambridge: Cambridge University Press.

Gentner, D., Rattermann, M. J., Markman, A. B., & Kotovsky, L. (1995). Two forces in the development of relational similarity. In T. J. Simon & G. S. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling* (pp.263-313). Hillsdale, NJ: Erlbaum.

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science, 10,* 277-300.

Halford, G. S. (1992). Analogical reasoning and conceptual complexity in cognitive development. *Human Development. Vol 35(4)*, (pp. 193-217).

Huttenlocher, J., Newcombe, N., & Sandberg, E. H. (1994). The coding of spatial location in young children. *Cognitive Psychology, 27,* 115-147.

Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in thedevelopment of relational similarity. *Child Development, 67,* (pp. 2797-2822).

Loewenstein, J., & Gentner, D. (1998). Relational language facilitates analogy in children. *Proceedings from the 20th Annual Conference of the Cognitive Science Society*, (pp. 615-619).

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

Uttal, D. H., & Schreiber, J. C., & DeLoache, J. S. (1995). Waiting to use a symbol: The effects of delay on children's use of models. *Child Development, 66,* (pp. 1875-1891).

Uttal, D. H., Gregg, V., Tan, L. S., Chamberlin, M., & Sines, A. (submitted). Connecting the dots: Children's use of a figure to facilitate mapping and search.