

Similarity Processing Depends on the Similarities Present: Effects of Relational Prominence in Similarity and Analogical Processing

Mark T. Keane (Mark.Keane@ucd.ie)

Department of Computer Science, Lukasiewicz Building,
University College Dublin, Belfield, Dublin 4 IRELAND

Deirdre Hackett (Deirdre.Hackett@erc.ie)

Educational Research Center, Dromcondra, Dublin 9 IRELAND

Jodi Davenport (Jodi@mit.edu)

Department of Brain & Cognitive Sciences, Massachusetts Institute of Technology
77 Massachusetts Avenue, Cambridge, MA 02139-4307 USA

Abstract

Several studies have shown that similarity judgements involve a process of structural alignment akin to analogical mapping. Some research has shown that performing a similarity judgement task produces more relational responding in a subsequent cross-mapping task, suggesting that similarity necessarily uses structural alignment. However, other research has shown that this effect disappears when procedural/material manipulations fail to emphasise the relational aspects of similar scenes. The present study confirms the latter findings showing that if relational similarities are less prominent in a material set then subjects respond in an object-based rather than a relational way. Importantly, these results show that similarity processing does not by necessity make use of structural alignment but that the similarity processing adopted is pluralistic and depends on properties of the presented materials.

Introduction

A considerable body of recent research has shown that similarity comparisons can involve a process of structural alignment (see e.g., Goldstone, 1994; Goldstone & Medin, 1994; Goldstone, Medin & Gentner, 1991; Markman & Gentner 1993a, 1993b, 1997; Medin, Goldstone & Gentner, 1993). Representationally, this view characterises knowledge as structured hierarchies encoding objects, object attributes, relations between objects and relations between relations. Given these representations it is assumed that similarity comparisons involve the alignment of relational structure to find the most structurally consistent match between two systems of concepts, that satisfies the constraints of parallel connectivity (if two relations match, their arguments

must match) and one-to-one mapping (that each item in one structure may only be mapped to one other item). Computationally, these ideas have been realised by a family of models that simulate analogical mapping (see e.g., Falkenhainer, Forbus & Gentner, 1989; Gentner, 1983, 1989; Holyoak & Thagard, 1989; Hummel & Holyoak, 1997; Keane, 1988, 1997; Keane & Brayshaw, 1988; Keane, Ledgeway & Duff, 1994; Veale & Keane, 1994, 1997, 1998). Indeed, structural alignment has been mooted as a unified account of a diverse range of phenomena including similarity, analogy, metaphor and concept combination (see Costello & Keane, 2000; Gentner, Holyoak & Kokinov, 2001; Keane & Costello, 2001).

Markman & Gentner (1993b) provided one of the key pieces of evidence supporting the role of structural alignment in similarity judgements. They used a one-shot mapping task in which subjects had to identify a cross-mapped object between two drawn scenes (see Appendix A). A cross-mapped object was defined as an object in one drawing that was perceptually similar to an object in a different relational role in the other drawing. So, for example, in the baseball scenes shown in Appendix A, the cross-mapped object would be the pitcher with the "C" on his uniform, because he is pitching in the upper scene and being pitched to in the other scene. Markman & Gentner have proposed that structural alignment is reflected in this task when subjects make relational responses (i.e., choosing the object in the same role) as opposed to object responses based on perceptual, feature similarity (i.e., choosing the perceptually similar object in a different role). The key manipulation asked participants to perform a similarity judgement task on the picture-pairs either before or after the mapping task. They found that

when participants made the similarity judgement *before* the mapping task they made more relational responses than when it was presented *after* the mapping task. Thus, the result strongly suggested that the similarity judgement task invoked a structural alignment process which then carried over to the mapping task increasing the proportion of relational responses (significantly, when an aesthetic-appreciation task was given before the mapping task no facilitation in relational responding was found).

Davenport & Keane (1999) queried these findings by pointing out that the materials and presentational procedure used may have contained unintended cues to promote relational responding. First, some of Markman & Gentner's materials had linguistic labels indicating the key relational similarity between the pictures (see e.g., the baseball pair in which "Pitch" is written). Second, the presentational procedure may have supported relational responding in that the all 8 stimuli presented were picture pairs with prominent relational similarities. Davenport & Keane found that when the linguistic cues were removed from the materials and the materials were mixed with fillers (involving simply object similarities) the similarity-judgement effect disappeared; that is, relational responding did not increase reliably when the similarity judgement task was made *before*, rather than *after* the mapping task. Significantly, Davenport & Keane also found that when the relational materials were blocked as a group and presented before the fillers there was increased relational responding relative to a condition in which the relational materials were randomly distributed among fillers in their order of presentation. This blocked-condition mimicked the presentational procedure used by Markman & Gentner. No reliable interaction was found between the ordering of the similarity judgement task (before or after mapping task) and the materials-order variable (blocked before or distributed among fillers). Hackett (2000) has subsequently replicated this materials order effect.

This pattern of results demands a very different account of similarity and mapping to that proposed by Markman & Gentner. First, it is no longer safe to assume that structural alignment is used in similarity judgements, as a matter of course, because relational responding does not follow prior similarity judgements. Second, it appears that relational responding is mainly dependent on the materials used and how they are presented. Specifically, the pattern of responding in the mapping task suggests that the *prominence of relational similarities* in the materials is the key variable affecting relational responding. This proposal best explains the evidence found:

- when linguistic cues are present that highlight key relational similarities then relational responding is seen (as in Markman & Gentner's original

materials)

- when several materials are consecutively presented with relational similarities then relational responding results (as in Markman & Gentner's study and Davenport & Keane's blocked condition)
- when materials with clear object similarities are mixed up with these relational materials then relational responding decreases (as in Davenport & Keane's distributed condition)

So, structural alignment is really only used when the prominence of relational similarities in the materials appear to demand it, perhaps with the default style being processing based on attribute similarities.

If this account is true then any manipulation that reduces the prominence of relational similarities in the materials should reduce relational responding. For instance, if we take the materials that have previously produced relational responding and add in additional object similarities then, on balance, the relational similarities should be less prominent. Hence, we should see reduced relational responding for these modified materials. For example, consider picture-pair A in Appendix II; the top picture shows a woman kicking a football with goal posts and a sun behind her and the bottom picture shows her being kicked by a child with some blocks and a clock in the background. This picture pair is quite sparse, a sparseness that lends a greater prominence to the kicking relation shown. Compare picture-pair A with picture-pair B in Appendix II; the latter also shows two kicking episodes but the scenes are much richer with more similar objects in the top and bottom pictures (houses, the sun, goal posts, etc). Although, both picture pairs have the same kicking incident, the greater frequency of object similarities in the richer pair should, if our hypothesis is correct, reduce the occurrence of relational responding relative to the sparser pair.

The present study examines this sparse versus rich manipulation, where the richer pairs were essentially the same scenes with added matching objects. We also attempted a further replication of the similarity-judgement effect by giving different groups a similarity judgement task either before or after the mapping task. As such, we had a 2 x 2 between-subject design where the variables were task-order (similarity task before or after mapping task) and material-type (sparse or rich materials). Following Davenport & Keane, all conditions presented the materials in a distributed fashion with the target materials being randomly distributed with fillers. We made two predictions in the study. First, that the similarity judgement task would have no effect on relational responding, confirming Davenport & Keane's finding. Second, that the sparse materials would produce more relational responding than the rich materials, as the prominence of the relational

similarities is reduced in the latter by the greater frequency of object similarities.

Method

Subjects. Forty-eight undergraduate students at University College Dublin took part voluntarily in the experiment and were randomly assigned to one of the four between-subjects conditions.

Stimuli. The stimuli for this experiment consisted of 8 pairs of pictures depicting causal scenes with matching relational structure and 16 filler pairs. Each of these 8 pairs contained a cross-mapping as operationalized by Markman and Gentner (1993b) in which a pair of perceptually-similar objects were shown which played different roles in the matching relational structure of the two scenes (see Appendix B for an example). In all four conditions, the 8 relational pictures were designed so that the perceptually-similar target object was in approximately the same spatial position in the picture pairs (e.g., the woman in the soccer materials is in the same central position in both pictures). Two versions of the relational materials were prepared. The rich set was created by adding similar objects to both scenes of the original sparse pictures used by Davenport & Keane (see Appendix II for an example). In all other respects, the picture-pairs were the same (e.g., in the placing of the arrow indicating the to-be-mapped object).

Eight of the filler pairs depicted comparable scenarios without matching relational structure (e.g. two beach scenes, one with a man surfing another with a child is building a sand castle) and the other 8 pairs did not match in either scenario or relational structure (e.g. a scene of an artist and a scene of a man in a grocery store). The fillers were the same as those used by Davenport & Keane.

The stimuli were presented in booklet form with one pair on each page (one picture above the other). The stimuli for the mapping task had an arrow placed above an object in the top scene. For the 8 target pairs this was the cross-mapped object, otherwise it was an object which appeared in both scenes. The stimuli used for the similarity rating task had a scale with the numbers 1 through 9 at the bottom of the page. The words Low Similarity appeared under the 1 and the words High Similarity appeared under the 9.

Booklets in the all conditions had a completely randomised presentation of the 24 pairs for both the mapping and the similarity tasks.

Procedure. As in Gentner & Markman's study, the first page of the mapping section of the booklet instructed subjects to draw a line from the object under the arrow to the object in the bottom scene that "best went with that object". The first page of the similarity judgement section instructed subjects to rate the similarity of the two scenes by circling a number

on the scale at the bottom of the page.

Subjects in the similarity-after conditions received a booklet with the mapping task followed by the similarity judgement task while subjects in the similarity-before condition received a booklet with the similarity judgement task first.

Subjects were tested in small groups of varying sizes and each experimental session took between 10 and 15 minutes.

Scoring. As in Gentner & Markman's study, participants' responses to the 8 target pairs in the mapping task were determined as an *object mapping* if a line was drawn from the cross-mapped object to a featurally-similar object in the bottom scene; a *relational mapping*, if a line was drawn to the object in the same relational role in the bottom scene; or a *spurious mapping* if a line was drawn to another, unrelated object. As in previous studies, spurious responses were removed prior to data analysis (5% of all responses made).




Figure 1: The Mean Proportion of Relational Responses in the Conditions of the Study

Results & Discussion

A two-way, between-subjects ANOVA found a reliable effect of material-type with a higher proportion of relational responding in the sparse conditions ($M = .37$) than in the rich conditions ($M=0.19$), though this difference is marginally reliable, $F(1,44) = 3.538, p=.06, MSE_{error} = 0.079$ (see Figure 1). Again, contrary to Markman and Gentner's predictions, no reliable effect was found for the task-order variable, $F(1,44) = 2.279, p > .10$. Indeed, the direction of the difference is in the opposite direction to that predicted; *less* relational responding occurs when the similarity judgement task is *before* as opposed to *after* the mapping task. Finally, there was no reliable interaction between task-order and material-type, $F(1,44) = 0.804, p > .10$.

General Discussion

There are two significant findings in this study. First, it replicates Davenport & Keane's finding that

relational responding is not influenced by a prior similarity judgement task; leaving open the question of whether similarity judgement involves structural alignment. Indeed, if we adopt the argument made by Markman & Gentner, we would be bound to conclude that structural alignment does not necessarily occur in similarity judgements. Second, and perhaps more surprisingly, it shows that structural alignment does not necessarily occur in the mapping task either. Rather, it appears that people sometimes respond on the basis of relational similarities and other times respond on the basis of object similarities. Furthermore, the key factor determining the mode of response lies in the nature of the materials themselves. When relational similarities are prominent in the given materials or across a set of consecutive materials then relational responding will result, but when these relational similarities are counter balanced by more object similarities in a given material or across a set of consecutive materials then object-based responding results.

Computationally, these findings present a number of challenges for models of similarity and analogy. They suggest that there are two distinct modes of processing for similarity judgement and mapping tasks. In one mode, relational correspondences are mainly used; this could be achieved by only computing relational matches (e.g., a type of selective attention to relations) or by computing both object and relational matches and then subsequently giving a greater weight to relational similarities. In the other mode, object attributes are mainly used: this could be achieved by only computing attribute matches (e.g., a type of selective attention to object features) or by computing both object and relational matches and then subsequently giving a greater weight to object similarities. Where you have two modes of processing there has to be a trigger for switching processing from one mode to the other. The empirical evidence suggests that this trigger is sensitive to the relative frequency of relational versus attribute similarities present in a stimulus pair and a set of stimulus pairs.

Do any current models have these sort of properties? Goldstone's (1994) interactive activation model gave a greater weight to attribute matches during early stages of processing with relational matches emerging later on; this model deals with the finding that under time pressures people rely more on attribute similarity (Medin & Goldstone, 1994). However, it is not immediately clear whether it would predict less relational responding in the rich versus the sparse materials. Furthermore, this model would need to have a history of previous similarity judgements to model the effects of consecutive relational materials. Another model, the MAX model (Goldstone, Medin & Gentner, 1991) pools relational and attribute similarities separately, with relational or attribute

responses being chosen based on the relative sizes of the two pools of similarities. This model might be able to deal with the relational prominence effect, as the rich materials might have a larger pool of attribute similarities relative to the sparse materials, leading to an object-based response. However, the MAX model does not have an account of the effects of consecutive materials as, again, it holds no history of previous similarity episodes. In short, no current model seems to be able to handle this evidence.

The current results leave open the question of whether similarities are selectively attended to with only a subset of all possible similarities being computed or whether all similarities are computed and then evaluated to decide on what response mode should be adopted. This is a key question to be settled by future empirical work. They also leave open the question of which mode of processing is the default mode or, in indeed, if there is a default mode. Attribute similarities seem to be computed more efficiently and easily (Goldstone & Medin, 1994). This mode of processing also has a laziness that is characteristic of human cognitive processing, when one considers people's mental sloth in making elaborative, bridging relational inferences.


However, there is one conclusion that is unavoidable given the current results: Namely, that comparison processes -- whether they be similarity or analogical processes -- are pluralistic rather than monolithic (see Medin, Goldstone & Gentner, 1994; Keane & Costello, 2001). This conclusion should shift the focus of research to the key issue of what variables influence the mode of processing adopted in this pluralistic computational environment.

References

- Costello, F. & Keane, M.T. (2000). Efficient creativity: Constraints on conceptual combination. *Cognitive Science*.
- Davenport, J. & Keane, M.T. (1999). Similarity & structural alignment: You can have one without the other. In *Proceedings of the Twenty-First Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Falkenhainer, B., Forbus, K.D., & Gentner, D. (1989). Structure-mapping engine. *Artificial Intelligence*, 41, 1-63.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. (1989). Mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*. Cambridge: CUP.
- Gentner, D., Holyoak, K.J., & Kokinov, B. (2001). *The Analogical Mind*. Cambridge, MASS; MIT Press.


- Goldstone, R.L. (1994). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Memory and Cognition*, 20, 3-28.
- Goldstone, R.L. & Medin, D.L. (1994). Time course of comparison. *Journal of Experimental Psychology: Language, Memory & Cognition*, 20, 29-50.
- Goldstone, R.L., Medin, D.L., & Gentner, D. (1991). Relational similarity and the non-independence of features in similarity judgments. *Cognitive Psychology*, 23, 222-262.
- Hackett, D. (2000). *Similarity judgements and the causes of relational responding*. MSc Thesis, Department of Computer Science, University College Dublin, Dublin, Ireland.
- Holyoak, K.J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295-355.
- Hummel, J.E. & Holyoak, K.J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104, 427-466.
- Keane, M.T. (1988). *Analogical Problem Solving*. Chichester, England: Ellis Horwood (Cognitive Science Series). [Simon & Schuster in N. America]
- Keane, M.T. (1997). What makes an analogy difficult ?: The effects of order and causal structure in analogical mapping. *Journal of Experimental Psychology: Language, Memory & Cognition*, 23, 946-967.
- Keane, M.T. & Costello, F. (2001). Why Conceptual Combination is Seldom Analogy. In D. Gentner, K.J. Holyoak, & B. Kokinov (Eds.), *The Analogical Mind*. Cambridge, MASS: MIT Press.
- Keane, M.T., & Brayshaw, M. (1988). The Incremental Analogical Machine: A computational model of analogy. In D. Sleeman (Ed.), *European Working Session on Machine Learning*. London: Pitman.
- Keane, M.T., Ledgeway, T. & Duff, S. (1994). Constraints on analogical mapping: A comparison of three models. *Cognitive Science*, 18, 387-438.
- Markman, A.B. & Gentner, D. (1993a). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language*, 32, 517-535.
- Markman, A.B. & Gentner, D. (1993b). Structural alignment during similarity comparisons. *Cognitive Psychology*, 25, 431-467.
- Markman, A.B. & Gentner, D. (1997). The effects of alignability on memory. *Psychological Science*, 5, 363-367.
- Medin, D.L., Goldstone, R.L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, 100, 254-278.
- Veale, T. & Keane, M.T. (1994). Belief modelling, intentionality and perlocution in metaphor comprehension. *Proceedings of the Sixteenth Annual Meeting of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Veale, T. & Keane, M.T. (1997). The competence of sub-optimal structure mapping on 'hard' analogies. *IJCAI'97: The 15th International Joint Conference on Artificial Intelligence*. Morgan Kaufmann.
- Veale, T. & Keane, M.T. (1998). Principle Differences in Structure Mapping. In K.J. Holyoak, D. Gentner, & B. Kokinov (Eds.), *Proceedings of Analogy '98*. New University of Bulgaria Press: Bulgaria.

Appendix I. The Baseball Materials Used by Markman & Gentner(1993b)



Appendix II. Examples of the materials used in the study showing the (a) spare kick material and (b) the rich kick material.

A.



B.

