

MODELING PERCEPTUAL SIMILARITY AS ANALOGY RESOLVES THE PARADOX OF DIFFERENCE DETECTION

Andrew Lovett¹

andrew-lovett@northwestern.edu

Eyal Sagi²

ermon@northwestern.edu

Qualitative Reasoning Group¹

Dedre Gentner²

gentner@northwestern.edu

Kenneth Forbus¹

forbus@northwestern.edu

Psychology Department²

Northwestern University, 2133 Sheridan Road
Evanston, IL 60208 USA

ABSTRACT

There is a paradoxical dissociation between recognizing *that* two stimuli are different and recognizing *how* they are different. We show that this dissociation can be captured by modeling perceptual similarity as a species of analogical processes. Using SME to model comparison, we show that the dissociation arises naturally from different stages in the analogical mapping process. Rather than relying on hand-coded input representations, our model uses an automatic, incremental encoding process to generate representations from the same stimuli as given to human participants.

INTRODUCTION

Although analogy was traditionally associated chiefly with creative discovery and problem-solving, there is increasing evidence that the same cognitive processes that humans utilize in abstract analogies may also be at work in concrete similarity comparisons (Markman & Gentner, 1996; Medin, Goldstone, & Gentner, 1993). The idea that a single process may underlie both literal similarity and analogy has garnered theoretical as well as empirical support (Gentner, 1983; Gentner & Markman, 1995; Goldstone & Medin, 1994).

One arena in which this approach has been particularly fruitful is the study of how people process differences. For example, Markman and Gentner (1996) found that when participants judged the similarity of pairs of

images, they were particularly sensitive to the *alignable differences* between the images: that is, to differences that correspond within the common structure (such as different objects in the same relational role). In contrast, they were less sensitive to *nonalignable* differences: differences that are not linked to the common structure or are linked in different ways. This finding is predicted by Gentner's (1983) structure-mapping theory of analogy, in which representations are compared by aligning their structure, thereby highlighting both the common structure and differences connected to it.

A further prediction is that because highly alignable pairs (with considerable common structure) are easier to compare than less alignable pairs, differences between them will be more easily noticed. This leads to the somewhat surprising prediction that people should be better able to identify differences between pairs of similar concepts than between pairs of dissimilar concepts—a prediction borne out in studies by Markman and Gentner, (1993) (see also Gentner & Markman, 1994; Gentner & Gunn, 2001). This finding also holds for perceptual images: People asked to list differences between pairs of images listed more differences for highly similar than for dissimilar images, despite the fact that the dissimilar images clearly had more potential differences between them (Markman & Gentner, 1996). The key principle here is that alignable differences are naturally salient. Therefore, when pairs are easily aligned (as with similar pairs), their differences “leap out”.

This finding that it is easier to say *how* two things are different for highly similar pairs seems at odds with a large body of work on the same-different task showing that people are faster to notice *that* stimuli are different for *dissimilar* pairs than for similar pairs (Farrell, 1985; Posner & Mitchell, 1967; Tversky, 1969). We suggest that the structure-mapping process offers a resolution of this seeming paradox (Markman & Gentner, 2005; Gentner & Sagi, 2006; Lovett et al., 2007).

Analogical comparison is modeled by the Structure-Mapping Engine (SME) (Falkenhainer, Forbus, & Gentner, 1989) as a multistep local-to-global process. It begins by finding local matches between identical elements (attributes and relations that exist in both representations). These local matches are coalesced into structurally consistent clusters (*kernels*), which are then merged to form a consistent global mapping. At this point further inferences can be drawn and alignable differences become salient.

Naming a specific difference between stimuli requires a full global mapping, and therefore depends on the alignability of the stimuli. However, some comparison tasks can be accomplished without the full process. Specifically, if two items are highly dissimilar, recognizing *that* they are different can often be done in the first (local matches) stage of processing. For example, if there aren't enough local matches to support an identity match, then a quick "different" response can be made without continuing the alignment process.

In previous work, we found evidence for the predicted dissociation using the same materials in both tasks (Gentner & Sagi, 2006). Given pairs of plants or pairs of heraldic shields, people were faster to say "different" for dissimilar pairs (e.g., the column pairs in Figure 1, A&C and B&D); but faster to name a difference for similar pairs (e.g., the row pairs, A&B and C&D). Finally, also as predicted, people were faster overall on the same-different task, which can sometimes be accomplished in the first stage of the process, than on the name-a-difference task. A follow-up simulation (Lovett, Sagi & Gentner, 2007)

demonstrated how this dissociation can be explained based on structure-mapping¹. SME was given simplified versions of the materials, automatically encoded using CogSketch (Forbus et al., 2008), a sketch understanding system (described later). SME found fewer local matches for dissimilar than for similar pairs—consistent with the finding of faster "different" RTs for dissimilar pairs—but computed stronger global matches for similar than for dissimilar pairs—consistent with the finding of faster name-a-difference RTs for similar pairs.




Figure 1: Sample stimuli from Gentner & Sagi, 2006. Images in the same row are high-sim pairs; images in the same column are low-sim pairs.

The Gentner and Sagi (2006) study and the Lovett et al. (2007) simulation demonstrated the predicted reversal: Same-different responses are faster for low-similarity pairs and name-a-difference responses are faster for high-similarity pairs. However, structure-mapping predicts an additional dissociation between the tasks. In the same-different task, participants' reaction times should be strongly affected by the similarity of the objects in the images being compared. If the objects are dissimilar, there will be few initial local matches, allowing a fast "different" RT (because participants

¹ The materials simulated were heraldic shields (Gentner & Sagi, 2006, Experiment 1).

can reject the possibility that the images are the same based on only the first stage of SME). However, object similarity should play a much weaker role in the name-a-difference task, because this task requires computing a complete global mapping. Thus the alignability of the relational structure of the images should be the major determinant of reaction time

This paper examines these predictions by independently varying object similarity and relational similarity. Before presenting the psychological study, we describe a computational model called PEC: Parallel Encoding and Comparison. PEC makes use of an automatic encoding system for perceptual images (CogSketch) that removes the need for hand-coded representations. We show that our model can be run on the same stimuli as were given to the human participants, and that it generates measures that correlate with human performance. In addition to testing the predictions of structure-mapping, our simulation permits us to test specific models of perceptual encoding like that embedded in PEC (Lovett et al., in press).

COMPUTATIONAL MODEL

In the PEC model, perceptual comparison is seen as an interaction between two separate, but interleaved, processes operating in parallel (Lovett et al., in press). The *encoding* process incrementally builds up representations of the two stimuli, beginning with low-level features like object shapes, and concluding with high-level relational structure. The *comparison* process (SME) computes an analogical mapping between the two representations, first finding local matches² and, over time, building up a globally coherent structural mapping. When these processes run to completion, the model can both determine whether the stimuli

are the same or different and name a specific difference between them. However, for very dissimilar stimuli, the model can quickly recognize that they are different based on finding a small number of local matches between the low-level features in the stimuli.

We begin by more fully describing the Structure-Mapping Engine, our model of comparison. We then describe our model of incremental encoding, which uses CogSketch to automatically generate representations. Finally, we show how these processes interact and how they can be used to generate different predictions for the two tasks.

Comparison: The Structure-Mapping Engine

The Structure-Mapping Engine (SME) (Falkenhainer, Forbus, & Gentner, 1989; Forbus & Oblinger, 1990) is a computational model of analogical comparison. It is based on Gentner's (1983) structure-mapping theory, according to which people compare two representations by aligning their common structure to compute a structurally consistent mapping. The alignment process is guided by structural consistency and by the *systematicity* principle: that people implicitly prefer to maximize the size and depth of the aligned structure.

SME takes as input two cases: a base and a target. Each case is a structural description made up of *entities*, *attributes* of entities, and *relations*. Lower-order relations hold between entities; higher-order relations hold between lower-order relations. SME computes mappings in three steps.

1: SME computes all possible local matches between statements in the base and target. Local matches must be either (a) identical attributes or relations; or (b) functions that are corresponding arguments of matched statements.

2: The local matches are coalesced into structurally consistent clusters (*kernels*)—partial mappings between the base and target.

3: Kernels that are consistent with each other—i.e., that do not violate structural consistency—are merged to form *global mappings*. A global

² As always in SME, the local match stage is not inherently bottom-up; matches are made at any level in parallel. However, in the case of incremental encoding, object matches may be discovered before relational matches simply because of the order in which these are encoded.

mapping is a maximally (or nearly maximally) large consistent mapping between the base and target.

SME computes one to three global mappings between a base and target, heuristically seeking to find mappings that maximize systematicity. Each mapping consists of a set of *correspondences* between elements and relations in the base and target; a *structural evaluation score*—a numerical similarity measure based on the size and systematicity of the mapping; and a list of *candidate inferences*—predicates connected to the common structure in the base but not initially present in the target. Importantly, reverse candidate inferences can also be computed. An *alignable difference* is detected when an inference and a reverse inference clash.

Encoding: CogSketch

We use CogSketch (Forbus et al., 2008) to automatically construct representations of visual stimuli. CogSketch is an open-domain, general-purpose sketch understanding system that constructs structural representations from human-drawn sketches and other line drawings. Unlike previous sketch understanding systems, which focus on recognizing the objects in the sketch, CogSketch focuses on capturing and interpreting the spatial relations among (and within) the entities, including perceptual and spatial organization.

By using CogSketch to construct the spatial representation given to SME, we can give the system the same PowerPoint materials given to human participants. This allows us to escape the well-known problem of hand-coding—that the model’s input representations may be (perhaps unknowingly) tailored to fit the model and the predictions. In addition, as we show in the next section, by explicitly modelling the encoding process, we can explore specific hypotheses concerning the time course of encoding.

Users create a sketch in CogSketch by drawing a set of objects, called *glyphs*. Given a set of glyphs, CogSketch automatically com-

putes a number of spatial relations between them. These include topological relations—such as whether two glyphs intersect or one lies within another—and positional relations, such as **right-of**. For compactness, positional relations are only computed between glyphs that CogSketch believes to be adjacent.

In addition to computing spatial relations between glyphs, CogSketch can also compare the shapes of glyphs. It does this by building up a structural representation of the edges of the glyphs and comparing them. CogSketch can thus recognize when two glyphs are the same shape. Typically, this information is used to create *shape attributes*, which are applied to every glyph of the same shape. For example, while CogSketch has no pre-existing shape knowledge, and thus no way of recognizing triangles, it can recognize that all the triangle glyphs in a sketch are the same shape and create a new attribute that is applied to them. Thus when CogSketch compares two sketches using SME, it will recognize that they contain glyphs with the same shapes.

We are also experimenting with an extension to CogSketch that identifies groups based on the Gestalt grouping principles of good continuation and proximity, as described below.

Incremental encoding

There is evidence from visual same-different tasks that participants do not encode all of a visual scene at one time. Rather, it appears that object attributes are encoded before relations (Sloutsky & Yarlus, in preparation). Some configural properties may also be encoded very early (Love, Rouder, and Wisniewski, 1999). For example, in Figure 2 from our study, we hypothesize that participants are likely to notice the row of three objects on the top before they recognize the individual objects that make up the image.

Based on this evidence, we have built a rough model of encoding in which perceptual information is made available incrementally, in three steps: (1) configural attributes are encoded; (2) object attributes are encoded; (3) relations are encoded (both relations between




Figure 2. Shape groupings may be noticed before individual shapes.

objects and higher-order relations between groups of objects)³. Configurative attributes are encoded by CogSketch for groups of objects that form a line (e.g., the top row in Figure 2) or pairs of objects that are adjacent (e.g., the bottom row).

Interaction between Encoding and Comparison

A further assumption in our PEC model of perceptual comparison is what we could term the “eager comparison” assumption: that the comparison process can begin its work before the encoding process is complete. That is, SME can use the partial results of the encoding process to begin matching. This means that at first, SME’s initial local match process will be operating chiefly over object attributes and local configurative attributes (such as the top row in Figure 2). This assumption of simultaneous encoding and matching processes is related to the interactive process proposed by Hofstadter and colleagues (French, 1995; Hofstadter & Mitchell, 1994).

In the PEC model, encoding and comparison interact as follows:

1: Configurative attributes for the two stimuli are given to SME, which finds local matches between them. If the number of local

matches is much lower than the number of elements in the base and target representations, the pair is judged as different.

2: Objects and their attributes are added to the representations, and SME looks for additional local matches in the updated representations. Again, if the number of local matches is relatively small, the pair is judged as different.

3: Relations are added to the stimulus representations. SME looks for additional local matches in the updated representations and then proceeds through its other stages, computing a global mapping between the representations. The candidate inferences of the global mapping are used to identify alignable differences. Only at this stage can the name-a-difference task be done⁴.

This model makes three predictions about human performance on the two tasks and how they are affected by object and relational similarity. First, overall performance should be faster on the same-different task than on the name-a-difference task. This is because participants can often recognize that stimuli are different based only on the number of local matches in the first stage of SME, but all three stages of SME are required to compute a global mapping and identify a specific difference. Second, in the same-different task, participants should be fast to say “different” if *either* object similarity or relational similarity is low. This is because either of these will result in a small number of local matches in the first stage of SME. Third, performance on the name-a-difference task should be faster for high relational similarity pairs. Object similarity should play a less significant role. This is because naming a difference depends on computing a structural mapping between the stimuli. Thus, overall, relational similarity should affect both tasks (though in opposite directions), while object similarity should only play

³ The evidence is inconclusive as to whether configurative or object attributes should be encoded first. Further, we suspect that higher-order relations may be encoded after first-order relations. For simplicity, the current PEC model encodes configurative attributes, object attributes, and relations in that order.

⁴ The name-a-difference task can of course be done without aligning the pairs, especially if there is a very salient difference (e.g., elephant vs. no elephant). However, all else being equal, people find it easier to notice alignable differences (Gentner & Markman, 1994).

a major role in the same-different task.

EXPERIMENT

We evaluated our model’s predictions in a psychological experiment in which we independently varied the object and relational similarity of the stimuli. The same stimuli were given to two sets of participants, with one set performing the same-different task and the other performing the name-a-difference task. In this way, we were able to measure the influence that object and relational similarity have on each of the similarity tasks.

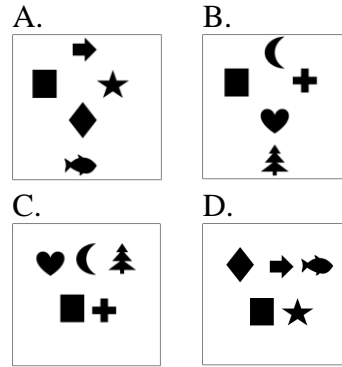
Participants and Materials

Fifty-three undergraduate students at Northwestern University participated, 20 in the same-different condition and 33 in the name-a-difference condition.

The materials were 60 images, each composed of 5 distinct objects (silhouettes) surrounded by a frame. Forty of the images (20 pairs) were designed such that in both images the spatial organization of the objects was highly similar (e.g., the rows in Figure 3). In half of these pairs (“high object similarity pairs”), 4 of the 5 objects were shared between the two images, while in the other half (“low object similarity pairs”), only 1 of the 5 objects was shared. The 20 pairs were then combined into groups of two pairs (e.g., Figure 3, A-D, G-J), such that the two pairs differed in their spatial organization but included the same objects. The remaining 20 images were used to create 20 pairs of identical images (“same” pairs).

Each participant saw 5 pairs from each of the experimental conditions (*high relational sim/high object sim*, *high relational sim/low object sim*, *low relational sim/high object sim*, and *low relational sim/low object sim*). In addition, participants in the same-different condition were also given the 20 “same” pairs. Finally, 10 pairs (five identical, five non-identical) consisting of arrangements of geometrical forms were used for training.

Low object similarity pairs



High object similarity pairs

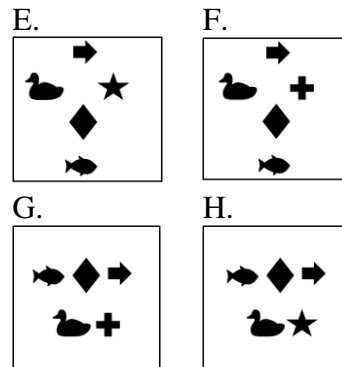


Figure 3. Sample stimuli. Within each set, images in the same row represent high relational similarity pairs; images in the same column represent low relational similarity pairs.

Procedure

The experiment was presented by computer. After completing a training phase, participants received the experimental pairs in two blocks of equal length. Each pair was preceded by a half-second fixation period during which a crosshair appeared at the center of the screen. The pair remained on the screen for 3000 ms.

In the same-different condition, participants judged whether the pair was identical or non-identical by pressing the left- or right-control key (counterbalanced). In the name-a-difference condition, participants typed in a difference between the two images. When the

	High Relational Similarity		Low Relational Similarity		
	Object Similarity	High	Low	High	Low
Same-Different		1.67 (.39)	1.24 (.25)	1.15 (.16)	1.11 (.13)
Name-a-Diff		3.93 (.50)	4.06 (.71)	4.27 (.84)	4.19 (.70)

Table 1. Mean response times by task and type of pair. Standard deviations are given in parenthesis.

participant responded (by making a same-different judgment or by starting to type a difference) or the 3000ms elapsed, the presented pair disappeared from the screen. In the name-a-difference condition, participants were then presented with a screen where they typed (or continued typing) the difference they had identified. For both tasks, the time between the onset of presentation of the pair and the response was recorded.

Results

Only correct “different” responses were used in the same-different analysis. This excluded approximately 9% of the “different” responses. Trials in which participants viewed different image pairs but responded “same” where removed (approximately 17% of the responses to different image pairs). The median response for each condition was then computed for each participant and each item. These medians provided the data points for the statistical analysis; their condition means are given in Table 1.

As predicted, same-different judgments were much faster than difference-identification (which took more than twice as long). Also as predicted, the two tasks showed different response patterns. In the same-different task, participants were faster to say “different” for pairs with low object similarity than pairs with high object similarity. “Different” responses were also faster for pairs with a substantially different spatial organization (low relational similarity) than for pairs that had a highly similar spatial organization. In contrast, participants in the name-a-difference condition were slower to identify a difference between low relational similarity pairs than between high relational similarity pairs. Their performance showed no effect of object similarity.

Repeated-measures ANOVA of Object Similarity x Relational Similarity for each task bore out these patterns. There was a significant effect of relational similarity in both tasks (though in opposite directions). (Same-Different: $F(1, 19) = 36.3$, $MS_e = .7$, $p < .01$; Name-a-Difference: $F(1, 32) = 7.2$, $MS_e = 4.3$, $p < .05$). However, object similarity only affected performance in the same-different task (Same-Different: $F(1, 19) = 48.1$, $MS_e = .7$, $p < .01$; Name-a-Difference: $F(1, 32) = .118$, $MS_e = 4.3$, *n.s.*). Likewise, the two variables showed a statistically significant interaction only when participants were performing the same-different task, but not for name-a-difference (Same-Different: $F(1, 19) = 20.8$, $MS_e = .7$, $p < .01$; Name-a-Difference: $F(1, 32) = 2.5$, $MS_e = 4.3$, *n.s.*).

Item ANOVAs for the two tasks showed similar patterns, revealing a main effect of relational similarity on both tasks (Same-Different: $F(1, 9) = 32.8$, $MS_e = .2$, $p < .01$; Name-a-Difference: $F(1, 9) = 20.3$, $MS_e = .3$, $p < .01$) but a main effect of object similarity only for participants in the same-different task (Same-Different: $F(1, 19) = 16.7$, $MS_e = .2$, $p < .01$; Name-a-Difference: $F(1, 32) = .11$, $MS_e = .3$, *n.s.*). As in the subject analysis, the interaction was significant only for participants in the same-different task (Same-Different: $F(1, 19) = 13.7$, $MS_e = .2$, $p < .01$; Name-a-Difference: $F(1, 32) = .63$, $MS_e = .3$, *n.s.*).

Finally, a preliminary analysis of the differences identified by participants suggests that participants were more likely to produce alignable differences when comparing images with high relational similarity than when comparing images with low relational similarity. For example, when comparing images E & F from Figure 3, 12 out of 17 participants (70%) produced alignable differences that contrasted the star in one image with the plus sign in the

other (e.g. “Right had star instead of a plus sign”). In contrast, when comparing images E & G, only 4 out of 14 participants (29%) produced alignable differences, and none of these identified the star in one image and the plus sign in the other. This pattern is similar to that observed by Markman and Gentner (1996) in which participants were more likely to identify differences that were alignable when two images were easier to align. We are currently conducting a more complete analysis of the identified differences.

Discussion

The results from the study matched our predictions. Participants were faster to recognize *that* the pairs were different than to name a particular difference. In the same-different task, participants were faster to recognize both low relational and low object similarity pairs as different, consistent with both local configural attributes and object attributes being encoded and compared quickly. In contrast, in the name-a-difference task, participants’ performance depended only on relational similarity, consistent with participants needing to perform a full structural mapping between the stimuli.

SIMULATION

We evaluated the PEC computational model by running it directly on the stimuli used with human participants. The questions of interest are (a) Would the model generate similarity measures that correlate with human reaction times on the tasks? And (b) Would the model show the same dissociations between the tasks as those found in the human data?

Procedure

We ran the PEC model on the same 40 pairings as were used in the psychological study: 20 high and 20 low relational similarity pairings, half high object similarity and half low object similarity.

The images were imported directly into CogSketch from PowerPoint. When PowerPoint images are imported, CogSketch automatically constructs a glyph for each PowerPoint shape⁵.

In evaluating our model, we looked at three measures generated by SME for each pairing. All measures were normalized based on the overall size of the two representations:

Local-config-matches: The number of local matches found by SME between representations in which only configural attributes have been encoded.

Local-object-matches: The number of local matches found by SME between representations in which only object attributes have been encoded.

Mapping-score: The structural evaluation score for global mappings computed by SME between the complete representations.

Our predictions for these measures were:

- (1) For the same-different task, the local-config-matches and local-object-matches measures should correlate with human performance, reflecting the claim that either a low number of configural matches or a low number of object matches should allow people to quickly determine that the images are different.
- (2) For the name-a-difference task, SME’s mapping-score should correlate with human performance, since structure-mapping predicts that stimuli that share more structure can be aligned more easily. However, object-matches should not correlate with performance.

Results

We consider each of the predictions in

⁵ Of the 10 object shapes (silhouettes) used in the psychological experiment, two were made up of multiple shapes in PowerPoint. These were simplified in PowerPoint so that CogSketch would build only one glyph for each of them. In addition, of the 40 images used, two were slightly touched up in PowerPoint because a shape was pasting into CogSketch badly. Other than these changes, the stimuli given to CogSketch were identical to the stimuli displayed to human participants.

turn. First, both local-config-matches and local-object-matches correlate with same-different performance (local-config-matches: $r=.64$, $p<.01$; local-object-matches: $r=.54$, $p<.01$). This positive correlation fits with the pattern that humans require more time to say “different” for more similar pairs (both at the object level and at the relational level).

Second, mapping-score correlates with name-a-difference performance ($r=-.47$, $p<.01$). However, local-object-matches does not correlate with name-a-difference ($r=-.25$, *n.s.*), producing our expected dissociation. This fits with the pattern that noticing a specific difference requires aligning the pair.

Discussion

As the results show, the PEC model is able to generate measures that correlate with human performance on both similarity tasks. Furthermore, our model’s measure of object similarity shows the expected dissociation between the tasks: it correlates with performance on the same-different task, in which object similarity matters, but not on the name-a-difference task, in which object similarity does not.

RELATED WORK

Several other researchers have used structural comparison to model human performance on visual similarity tasks, including the same-different task (Goldstone & Medin, 1994) and similarity rating tasks (Larkey & Markman, 2005; Taylor & Hummel, 2007). However, to our knowledge none of these models have been evaluated on the name-a-difference task.

CONCLUSIONS & FUTURE WORK

Structure mapping explains the reversal between the same-different and name-a-different tasks, wherein it is easier to see *that* dissimilar stimuli are different, but it is easier to say *how* similar stimuli are different. It also explains why different types of similarity play

distinct roles in each task. In the same-different task, both object similarity and relational similarity affect how quickly people can recognize that they are different. Both the configural information and the attributes of objects are encoded and compared quickly. If there are relatively few local matches in the first stage of comparison—insufficient for an identify match—then the pair can be immediately classified as different.

In contrast, in the name-a-difference task, the relational similarity of the pair plays a major role, while the attributes of objects are much less important. Images that share the same relational structure are much easier to align via structure-mapping.

One remaining question is whether our model can correctly predict the specific differences identified by participants in the name-a-difference task. As discussed earlier, we are currently analyzing these differences. We plan to test whether the differences stated by participants match the alignable differences computed by SME in our simulation.

ACKNOWLEDGEMENTS

This work was supported by NSF SLC Grant SBE-0541957, the Spatial Intelligence and Learning Center (SILC).

REFERENCES

- Falkenhainer, B., Forbus, K. D., Gentner, D. (1989). The Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence*, 41, 1-63.
- Farell, B. (1985). Same-Difference Judgments: A Review of Current Controversies in Perceptual Comparisons. *Psychological Bulletin*, 98, 419-456.
- Forbus, K. & Oblinger, D. (1990). Making SME Greedy and Pragmatic. In: *Proceedings of the 12th Annual Meeting of the Cognitive Science Society*.
- Forbus, K., Usher, J., Lovett, A., Lockwood, K., Wetzell, J. (2008). CogSketch: Open-Domain Sketch Understanding for Cogni-

- tive Science Research and for Education. In: Proceedings of the Fifth Eurographics Workshop on SBIM.
- French, R. M. (1995). *The Subtlety of Sameness: A Theory and Computer Model of Analogy-Making*. Cambridge, MA: The MIT Press.
- Gentner, D. (1983). Structure-Mapping: A Theoretical Framework for Analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. & Gunn, V. (2001). Structural Alignment Facilitates the Noticing of Differences. *Memory and Cognition*, 29, 565-577.
- Gentner, D. & Markman, A. B. (1995). Similarity is Like Analogy: Structural Alignment in Comparison. In: Cacciari, C. (ed.), *Similarity in Language, Thought and Perception*. Brussels: BREPOLs.
- Gentner, D. & Markman, A. B. (1994). Structural Alignment in Comparison: No Difference without Similarity. *Psychological Science*, 5, 152-158.
- Gentner, D. & Sagi, E. (2006). Does "Different" Imply a Difference? A Comparison of Two Tasks. In: Proceedings of the 28th Annual Meeting of the Cognitive Science Society.
- Goldstone, R. L. & Medin, D. L. (1994). Time Course of Comparison. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 29-50.
- Gentner, D. & Medina, J. (1998). Similarity and the Development of Rules. *Cognition*, 65, 263-297.
- Hofstadter, D. R. & Mitchell, M. (1994). The Copycat Project: A Model of Mental Fluidity and Analogy-Making. In Holyoak, K. J. & Barnden, J. A. (eds.), *Advances in Connectionist and Neural Computation Theory: Vol. 2. Analogical Connections*. Norwood, NJ: Ablex..
- Larkey, L. B. & Markman, A. B. (2005). Processes of Similarity Judgment. *Cognitive Science*, 29, 1061-1076.
- Love, B. C., Rouder, J. N., Wisniewski, E. J. (1999). A Structural Account of Global and Local Processing. *Cognitive Psychology*, 38, 291-316.
- Lovett, A., Gentner, D., Forbus, K., Sagi, E. (in press). Using Analogical Mapping to Simulate Time-Course Phenomena in Perceptual Similarity. *Cognitive Systems Research, Special Issue on Analogies – Integrating Cognitive Abilities*.
- Lovett, A., Sagi, E., Gentner, D. (2007). Analogy as a Mechanism for Comparison. In: Proceedings of Analogies: Integrating Multiple Cognitive Abilities.
- Markman, A. B. & Gentner, D. (2005). Nonintentional Similarity Processing. In: Hassin, R., Bargh, J.A., Uleman, J.S. (eds.) *The New Unconscious*, New York: Oxford University Press.
- Markman, A. B. & Gentner, D. (1996). Commonalities and Differences in Similarity Comparisons. *Memory & Cognition*, 24(2), 235-249.
- Markman, A. B. & Gentner, D. (1993). Splitting the Differences: A Structural Alignment View of Similarity. *Journal of Memory and Language*, 32, 517-535.
- Medin, D. L., Goldstone, R. L., Gentner, D. (1993). Respects for Similarity. *Psychological Review*, 100, 254-278.
- Posner, M. I., & Mitchell, R. F. (1967) Chronometric Analysis of Classification. *Psychological Review*, 74, 392-409.
- Sloutsky, V. M. & Yarlas, A. S. (in preparation). Processing of Information Structure: Mental Representations of Elements and Relations.
- Taylor, E. G. & Hummel, J. E. (2007). Perspectives on Similarity from the LISA Model. In: Proceedings of Analogies: Integrating Multiple Cognitive Abilities.
- Tversky, B. (1969). Pictorial and Verbal Encoding in a Short-Term Memory Task. *Perception & Psychophysics*, 6, 225-233.