Title

Similarity Between Propositional Elements Does Not Always Determine Judgments of Analogical Relatedness

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 30(30)

ISSN

1069-7977

Authors

Minervino, Ricardo A.
Oberholzer, Nicolas
Trench, Maximo

Publication Date

2008

Peer reviewed

Similarity between Propositional Elements does not always determine Judgments of Analogical Relatedness

Ricardo A. Minervino (minervinor@jetband.com.ar)
Departamento de Psicología, Facultad de Ciencias de la Educación, Universidad Nacional del Comahue, H. Irigoyen 2000, 8324 Cipolletti, Río Negro, Argentina

Nicolás Oberholzer (oberholzer@fibertel.com.ar)
Facultad de Psicología, Universidad de Buenos Aires, Independencia 3065, Piso 3
Buenos Aires, 1425 Argentina

Máximo Trench (maxtrench@fibertel.com.ar)
Facultad de Psicología, Universidad de Buenos Aires, Independencia 3065, Piso 3
Buenos Aires, 1425 Argentina

Abstract

Most computational models of analogical mapping and evaluation of analogical relatedness (e.g., SME, ACME and LISA) reduce the effects of semantics on these processes to the influence of similarities between propositional elements to be paired. Two experiments were carried out to show that people do not always follow this kind of similarities. In Experiment 1, when comparing a single proposition base analog with two alternative target analogs, participants judged as more analogous those that did not share that type of similarity. In Experiment 2 participants solved ambiguous mappings between propositions framed within systems of relations and tasks of cause identification. They favored matchings between propositions lacking element to element similarity. The implications of these results for computational models of analogical thinking are discussed.

Keywords: analogy; similarity; mapping; rerepresentation.

Semantics in the dominant theories of analogical mapping

Analogy is a powerful mechanism that takes part in many and diverse cognitive tasks (Gentner, Holyoak & Kokinov, 2001). Through a mapping between a well-understood *base analog* (BA) and a less-understood *target analog* (TA), inferences can be drawn to enhance the representation of the TA (Gentner, 1983; Holyoak & Thagard, 1995). The structure-mapping theory (Gentner, 1983, 1989) and the multiconstraint theory (Holyoak & Thagard, 1989; Hummel & Holyoak, 1997) have dominated the discussion of analogical mapping and inference generation.

The structure-mapping theory postulates that knowledge is represented in propositional form, and distinguishes between: (a) *entities*: single elements that stand for objects; (b) *attributes*: unary predicates representing properties of objects; (c) *first-order relations*: multiplace predicates that link two or more objects; (d) *higher order relations*: predicates that link relations themselves. According to this theory, in an analogy the two situations being compared share a relational structure (systems of relations governed by higher order relations), despite existing differences in the attributes of the objects that compose the situations. In literal similarity, both relational structures and attributes are shared (Gentner, 1989). SME (Falkenhainer, Forbus & Gentner, 1989), the computational implementation of the theory, is a symbolic system which takes as inputs propositional descriptions of the BA and the TA, and finds the maximal (i.e., largest and deepest) coherent relational match between the two, leaving aside isolated relations and not considering attributes. Two elements are allowed to be mapped only if they satisfy the following initial conditions: (a) *formal identity*: elements must be of the same formal type (objects, n-place relations, etc.); (b) *semantic identity for relations*: relations can be mapped only if they are identical in meaning. Once all local matches are generated, the program incrementally coalesces them into a few global mappings. Such mappings are *structurally consistent*, that is, they satisfy the following constraints: (a) *parallel connectivity*: if two predicates are put in correspondence, their arguments must also be mapped; (b) *one-to-one mapping*: each element in the BA must map to at most one element in the TA and vice-versa. SME uses the established mappings to suggest hypothesis about the TA. Finally, each global mapping is given a syntactic evaluation based on the number of local matches and the depth of the system of matches (the *systematicity principle*). To make the semantic identity condition for relations more flexible, Yan, Forbus and Gentner (2003) coupled SME with rerepresentation mechanisms aimed at detecting underlying commonalities between similar but non-identical relations.

The multiconstraint theory conceives mapping as determined by a conjunction of syntactic, semantic and pragmatic constraints. ACME (Holyoak & Thagard, 1989), the first computational implementation of the theory, is a hybrid system that combines propositional representations with connectionist-style processing. The program builds a network in which nodes stand for mapping hypotheses between formally identical elements, and weighted links between nodes represent constraints. Symmetric excitatory links are created between mapping hypotheses that satisfy the constraint of parallel connectivity, and symmetric inhibitory links are generated between mapping hypotheses that violate the one-to-one constraint. Pragmatic constraints are instantiated in ACME through a pragmatic unit that favors mapping hypotheses involving elements that are deemed relevant for the goals of the analogist. The program implements the constraint of semantic similarity by connecting mapping hypotheses to a special semantic unit from which they receive an activation proportional to the degree of previously known similarity between elements (these similarity scores are provided by the programmers following criteria akin to those entailed in IS-A networks). Once the connectionist mapping-network is constructed, an iterative algorithm operates to update connection weights. The mapping network settles as a result of having identified a set of correspondences that represent an optimal mapping. Albeit implicitly, ACME implements an evaluation criterion very similar to SME's systematicity principle, favoring bigger and deeper mappings. ACME's treatment of semantic similarity differs from SME's in two important aspects: (a) ACME takes into account similarities between objects during mapping—according to this theory these similarities have an effect on the evaluation of analogical relatedness, while in the structure-mapping theory they only count in judgments of literal similarity—; (b) if two propositions maintain identical formal positions within their analogs, they can be mapped even if their relations are not semantically similar—in this way ACME is supposedly able to discover new similarities (Hoyoak & Thagard, 1995).

LISA (Hummel & Holyoak, 1997, 2003), the second computational implementation of the multiconstraint theory, is a hybrid system that combines the semantic flexibility of connectionist architectures with the sensitivity to structure provided by symbolic models. Case roles and objects are represented in working memory (WM) as distributed patterns of activation over a collection of semantic units. When a proposition such as LOVE (bill, mary) gets activated, the semantic primitives of *lover* (e.g., *emotion* and *strong*), fire in synchrony with the semantic primitives of *Bill* (e.g., *male* and *adult*), while units representing the beloved role are synchronized with units representing *Mary*. Bindings are encoded and retrieved from long-term memory (LTM) through a hierarchy of structure units. At the bottom, predicate and object units link objects and case roles to their semantic primitives. One level above, sub-proposition units (SP) bind case roles to their fillers. At the top, proposition units (P) bind together two or more SPs. In LISA's LTM, two analogs are represented as non-overlapping sets of structure units. When the semantic primitives of a TA fire in WM, structure units at all levels compete in responding to patterns in WM. Across alternative base analogs, structure units of the same type are linked by inhibitory connections. Within a base analog, structure units of different type are linked by excitatory connections, leading elements in base and target to fire in synchrony. This synchrony of firing gets preserved in LTM through static bindings called *mapping connections*. Unlike ACME, pragmatic constraints don't need to be explicitly represented. As LISA's performance is sensitive to order effects, the programmers can allow "important" elements to be entered first. Due to its distributed representations, LISA doesn't need to be externally provided with similarity scores, since overlapping semantic units represent similarity between concepts.

Shortcomings of the dominant theories of mapping in their treatment of semantics

Most theories of analogical reasoning accept that two situations can be analogous even when their corresponding elements are not initially represented as having identical meaning (Gentner & Kurtz, 2006). They require, however, that some kind of identity between initially similar (but non-identical) elements could be found, and they propose several mechanisms for identifying those identities. We argue that a limitation underlying all these mechanisms consists in comparing base and target propositional elements separately from the analogs and analogies in which they take part (for similar criticisms, see Falkenhainer, 1990; Hofstadter & FARG, 1995). Sometimes, this type of isolated comparisons could turn out to be misleading. Consider, for example, the following BA: *Peter tried to unlock the house with the car remote key*. Suppose now that you are confronted with two TAs: (a) *Peter attempted to uncork the bottle with the motorcycle key* (TA1); (b) *Peter attempted to wash his hands with the hair gel* (TA2). Which would you consider as more analogous to the BA? Despite the existing similarities between corresponding elements of the BA and the TA1, some people will probably consider the TA2 to be more analogous to the BA, considering them to be two cases of, say, "being absent-minded and making action slips". This descriptor seems to be less trivial than the one that could be derived by simply concatenating the supraordinate concepts induced from mapped propositional elements of BA and TA1 (i.e., "trying to open something with a vehicle key").

Confronted with the task above, ACME would prefer the mapping between the BA and the TA1, reflecting the higher degree of similarity between their corresponding relations (UNLOCK↔UNCORK) and objects (*house↔bottle*, *car remote key↔motorcycle key*). The alternative mapping implies matching less similar relations (UNLOCK↔WASH) and objects (*house↔hands, car remote key↔hair gel*).

LISA would resolve this ambiguous mapping in a similar way. When the SP for Peter+unlocker fires in LISA's WM, it activates units for *Peter* and *unlocker*, transferring top-down activation to their semantic units. (e.g., *human*, *male*, *adult* and *open1*, *unlock1*, *physical1*, respectively). This pattern of semantic units in WM will excite all levels of structure units corresponding to the TAs, which will compete for activation. *Human*, *male* and *adult* will evenly activate both units for *Peter* within the targets (i.e., the unit for *Peter* as bound to *uncorker* and the unit for *Peter* as bound to *washer*). At the same time, *open1*, *physical1* will excite the predicate unit *uncorker*, but only *physical1* will excite *washer*. As *uncork1 and open1* will inhibit *washer*, LISA begins to act as if *unlocker* corresponded to *uncorker*. Analogous operations will cause LISA to act as if *house* corresponded to *bottle* and *unlocked* corresponded to *uncorked*. Accordingly, triggered by the SP car remote key+instrument that was built for the BA, LISA will map *car remote key* to *motorcycle key* and *instrument* to *instrument*. Because *uncorker* is more active than *washer*, the SP Peter+uncorker will receive more bottom-up input than the SP Peter+washer. Analogous support will be obtained by the target SP units bottle+uncorked and motorcycle-key+instrument. Finally, as SPs excite the P units to which they belong, the unit for UNCORK (peter, bottle, motorcycle key) will become more active than WASH (peter, hand, hair gel). LISA will thus prefer mapping UNLOCK (peter, house, car remote key) to UNCORK (peter, bottle, motorcycle key).

Even though SME would only take into account semantic similarities between relations but not between attributes, it would still choose the TA1 as more analogous. In fact, the program would need to apply the rerepresentation mechanisms proposed by Yan, Forbus and Gentner (2003). Possibly the most suitable one for the example we are analyzing would be *decomposition*. This mechanism involves breaking down predicate representations into the subcomponents that encode the meaning of the relational term. Such decompositions are then compared to discover identity matches among the components. In BA-TA1 applying decomposition could result in:

BA: UNLOCK (house) → CAUSE (UNLOCK (house), OPEN (house))

TA1: UNCORK (bottle) → CAUSE (UNCORK (bottle), OPEN (bottle)).

In this case an identical semantic element (*open*) is found within the meanings of both predicates. However, it is difficult to imagine how this method would reveal any identity between the BA and the TA2, since *unlock* and *wash* do not seem to share any identical semantic elements. Hence, SME augmented with these rerepresentational methods would also map the BA to the TA1.

It is clear that while these programs could be able to generate a common descriptor for BA-TA1, such as "open something with a vehicle key", they cannot identify the commonalities that give rise to a descriptor such as "being absent minded and making action slips". As a consequence, they are not able to map the BA to the TA2. What leads all these programs to choose the TA1 is the fact that they are considering similarities between corresponding base and target propositional elements in isolation. While ACME estimates this type of similarity according to the closeness of two elements within IS-A hierarchies, in LISA this similarity is implicitly represented by the co-activation of semantic primitives in WM. In SME, this similarity can be revealed via re-representation mechanisms.

We will use the term *element similarities* to refer to similarity between propositional elements. As the example above suggests, element similarities are not always sufficient to find the best analogical mapping, and can even be misleading if superficial pairings that compete with more profound ones are favored. Our first experiment was developed to show that two analogs that share element similarities could frequently be considered less analogous than two analogs that do not share them. In other words, that people do not always behave like SME, ACME and LISA.

Experiment 1

Our first experiment followed Gentner and Kurtz's (2006, Exp. 1) procedure. For each task, participants received a BA and two TAs. As these analogs were formally identical, a one-to-one mapping could be built between the BA and any of the two TAs. Participants were asked to choose the TA they considered more analogous to the BA. The TA1 was built substituting one or two verbs and one or two nouns in the BA with verbs and nouns with very similar meaning (e.g. *uncork* for *unlock*, *motorcycle key* for *car remote key*). In contrast, in the TA2 the same base verbs and nouns were replaced by verbs and nouns with less or no similar meaning (e.g. *wash* for *unlock*, *hair gel* for *car remote key*). The BA and the TA1 allowed a common description that could be derived by concatenating the supraordinate concepts of the mapped elements (e.g., *to open something with a vehicle key*). The BA and the TA2 allowed a common description not derivable from element similarities, but that was in our intuition less trivial[1] than the former (e.g., *being absent minded and making action slips*). See Table 1 for an example.

Table 1: Example of experimental materials.

Analog	Example
BA	John *gave*[a] *sweets* to a girl in his class.
TA1	John *shared* a *chocolate* with a girl in his class.
TA2	John *wrote* a *poem* for a girl in his class.

[a] In the original Spanish version "regalar" (i.e., to give a present).

As was analyzed, in a task with this structure, SME, ACME and LISA would map the units maintaining higher element similarities. We predicted that participants would prefer instead the mapping between units that involved lower element similarities but admitted a less trivial descriptor. In order to have independent measures of the degree of element similarities between propositional elements to be matched, we asked an independent group to rate the similarity between the BA elements and the corresponding elements in the TA1 and the TA2.

Method

Participants Sixty undergraduate students of Psychology at University of Buenos Aires took part in the experiment. Thirty received the analogical relatedness evaluation task (analogy group) and 30 the element similarity rating task (similarity rating group).

Materials Participants in the analogy group received six critical and six filler analogical tasks. Fillers were constructed in the same way as critical materials. To prevent participants from inducing an association between non-trivial descriptors and lack of elements similarities, one of

[1] The kind of descriptors applicable to a pair of analogs (e.g., trivial vs. non trivial) was not an independent variable in our experiment, but just an informal intervention to prevent participants from following element similarities.

the TA fillers shared both element similarities and the non-trivial descriptor with the BA, and the other neither of them. The order of presentation of the TAs and the 12 tasks were counterbalanced. The similarity-rating group received a list with 32 pairs of concepts to be evaluated using a 5-point Likert scale ranging from 1 (no similarity) to 5 (high similarity). The order of presentation of the pairs was counterbalanced.

Design and Procedure The independent variable was the degree of element similarities between base and target elements (high vs. low), a within-subjects variable. The dependent variable was the chosen TA. For each of the tasks, participants in the analogy group had to read a first analog, consisting in a sentence that described a simple fact. Afterwards, they were presented with two other analogs, and they were asked "Which of these two facts do you find more analogous to this first one?" The similarity-rating group was asked to evaluate to what extent they considered each pair of elements similar.

Results and Discussion

The similarity rating task yielded 16 critical results, each of which came out of comparing the rating mean for more similar pairing against the rating mean for the corresponding less similar pairing (e.g., rating mean for *give-share* vs. rating mean for *give-write*). We computed t statistics for these comparisons. In 13 cases the mean of the similarity ratings for the more similar pairings was greater than the mean for the less similar pairings. The three remaining comparisons showed non-significant differences (due to space restrictions, we cannot give more details about these results).

Table 2 shows the percentages of the analogy group participants that chose, for each of the critical tasks, TA1 or TA2. Data show that even though propositional elements of the BA were rated as more similar to their corresponding elements in the TA1 than in the TA2, participants in the analogy group chose the TA2 as more analogous to BA than TA1 in five of the six tasks (we found no trend in the remaining task). Thus, participants' responses to the analogical relatedness task diverged from the mapping that would be chosen by programs like SME, ACME or LISA. People passed over element-to-element similarities, and seemed to favor general descriptors that, although not derivable via supraordination of corresponding propositional elements, were nevertheless applicable to propositions as wholes.

Table 2: Percentages of TAs choices in the analogy group.

T	TA1	TA2	χ^2 (1, N =30)
1	30%	70%	9.60^b
2	20%	80%	21.60^b
3	47%	53%	.27
4	33%	67%	6.67^b
5	27%	73%	13.07^b
6	10%	90%	38.40^b

Note. T: Task; [b] Significant at .01 level

As Gentner and Kurtz's (2006) stimuli, our analogs consisted in first order propositions. However, the concept of analogy refers, *strictu sensu,* to a comparison between two systems of relations (Gentner, 1989; Holyoak & Thagard, 1995). In addition, the employed analogical relatedness task is removed from the type of purposeful activities in which analogical reasoning routinely participates (e.g., problem solving, argumentation, etc.). In Experiment 2 we sought to extend our findings using more natural tasks, as well as materials that comprise mappings between systems of relations.

Experiment 2

In Experiment 2 participants used analogical reasoning to identify the cause of a target effect. They were given a BA consisting in a base cause (BC) and its effect. Afterwards, they were told that the base effect reoccurred as a consequence of a cause *analogous* to BC, and were given two alternative target causes (TC1 and TC2) from which to choose. In this way, the task used in Experiment 1 —judging which of two TAs was more analogous to a BA— shifted to one of deciding an ambiguous mapping between a base proposition and two alternative target propositions (both arguments of a second-order causal relation). Adapting the example given in Experiment 1, suppose that a BA stated that *Peter tried to open the house with the car's remote key* (BC), causing Mary, his wife, to get worried (effect). The TA would state that, at a later time, *Mary got worried again*. Then, participants would be presented with two events: (a) *Peter attempted to uncork the bottle with the motorcycle key* (TC1); (b) *Peter attempted to wash his hands with the hair gel* (TC2). Under the supposition that what caused the latter worry was *analogous* to what caused the former, participants would have to choose between TC1 and TC2.

The BC could be matched with any of the two TCs under formal and pragmatic considerations. However, the propositional elements of the BC shared more element similarities with the corresponding elements in the TC1 than in the TC2. We asked an independent group of participants to rate the similarity between the BA elements and the corresponding elements in the TA1 and the TA2. Table 3 shows an example of the employed materials.

Table 3: Example of experimental materials.

Item	Example
BC	Lucas decided to *go* to *spinning classes*
B effect	Lucas' wife became really happy
TC1	Lucas decided to *attend relaxation lessons*
TC2	Lucas decided to *give up fast food*
T effect	Lucas' wife became really happy again

As already analyzed, in a mapping task like this, SME, ACME and LISA would prefer mapping the two causes whose corresponding elements maintain higher element similarities. We predicted that participants in this group would prefer instead the mapping that involved lower element similarities, but enabled, as we supposed, the generation of more meaningful descriptors. For each analogical task, participants in a control group received the TA but not the BA, and were asked to guess which of the two TCs could have generated the target effect. If target-only participants showed no bias towards a given TC, then preferences among experimental participants could be attributed to the influence of the BA, and not to a higher intrinsic plausibility of a TC as a cause for the target effect.

Method

Participants Ninety students of Psychology at University of Buenos Aires took part in the experiment. Thirty were randomly assigned to the analogy group, 30 to the target-only group and 30 to the similarity-rating group.

Materials Participants were presented with six critical tasks and six filler tasks. The TCs were built in the same way as TAs in Experiment 1, but were coupled with a common effect. In both groups, presentation of the TCs was followed by a question asking to choose one of them. The order of the tasks and the TCs was counterbalanced.

Design and Procedure The independent variables were similarities between elements of the BC and the TCs (high vs. low), a within-subjects variable, and condition (analogy vs. target-only), a between-groups variable. The dependent variable was the chosen TC. Participants in the analogy group read the text describing the BA. Afterwards, they read the text describing the target effect. Once they had finished reading the BA and the target effect, participants were asked: "Assuming that what caused the first event was *analogous* to what caused this last event, which of these two causes would you choose?" Participants in the target-only group were told that they would be given two unrelated tasks, one of test comprehension and one of identification of causes. After reading the irrelevant text, they were presented with the target and the TCs. They were asked to choose, following their own criteria, the TC they considered to be more plausible as a cause for that event.

Results and Discussion

The element-to-element similarity rating task yielded 18 critical results. We computed *t* statistics for these comparisons. In 16 cases the means of the similarity ratings for the more similar pairings were greater than the means for the less similar pairings. The two remaining comparisons showed non-significant differences. Table 4 shows the percentages of participants that chose TC1 and TC2 for the analogy and target-only groups, together with their Chi square statistics.

Table 4: Percentages of TCs choices in the analogy group and the target only group.

	Analogy group			Target-only group			A
T	TC1	TC2	χ^2	TC1	TC2	χ^2	χ^2
1	33%	67%	6.67[b]	60%	40%	2.40	4.29[a]
2	27%	73%	13.07[b]	57%	43%	1.07	5.55[a]
3	23%	77%	17.07[b]	50%	50%	.00	4.27[a]
4	43%	57%	1.07	67%	33%	6.67[b]	3.30
5	17%	83%	26.67[b]	63%	37%	4.27[a]	13.61[b]
6	23%	77%	17.07[b]	53%	47%	.27	5.71[a]

Note. In all cases $N = 30$ and $df = 1$. A: Association between choice and condition; T: Task. [a] Significant at .05 level [b] Significant at .01 level.

In five critical tasks participants preferred TC2 to TC1. Data thus replicate results obtained in Experiment 1 (we found no trend in the remaining task), with the upgrading that the compared propositions of the first experiment were now framed within a natural analogical task. People solved the ambiguous mappings differently from SME, ACME and LISA. The preference for TC2 in the analogy group cannot be attributed, in any of the five tasks, to a higher intrinsic plausibility of TC2 in the TA, since the target-only group showed a preference for the TC1 or no preference at all. In five out of the six tasks, we found an association between TC choice and condition.

General Discussion

Several authors (e.g., Falkenhainer, 1990; Hofstadter & the FARG Group, 1995) have argued that both SME and ACME inconveniently reduce the influence of semantics on mapping to predetermined general-purpose similarities between propositional elements of the analogs. We extended this criticism to LISA.

Experiment 1 demonstrated that when judging analogical relatedness between first order propositions analogs people pass over element similarities and seem to follow global descriptors that cannot be derived from element similarities. Experiment 2 replicated these results inserting first order propositions within systems of relations, and replacing the analogical relatedness task with a more natural task involving identification of causes. Unlike SME, ACME and LISA, which would be biased towards element similarities, our participants proved to resist them in favor of non trivial descriptors that were not derivable from concatenating supraordinated concepts of the paired elements.

The rerepresentation mechanisms that have been proposed to discover identities between initially dissimilar facts fall short of accounting for people's ability to generate descriptors that do not emerge out of element similarities. In an interesting example provided by Kurtz (2005, p. 449) illustrating human flexibility in making analogies, the BA *Amy wants to date Bill because he wears a leather jacket* was paired to the TA *Richard admires Michelle because she drives a Saab 900 convertible*. According to Kurtz, *wanting to date* and *admiring* are two cases of "a positive stance and forward attitude towards a person", at the time that *wears* and *drives* "connote ownership of an object *in a manner that contributes to personal identity*" (italics added). Its apparent to us that this type of similarities cannot be identified through IS-A networks, decompositions, or coactivations of semantic units within distributed representations. They seem to emerge from comparing whole facts via mechanisms that deserve further exploration. Our data suggest that the development of rerepresentation mechanisms capable of opportunistically compensating our general bias towards element similarities —well captured by dominant computational models and supported by empirical evidence (cf., e.g., Gentner & Kurtz, 2006; Reeves & Weisberg, 1994)—should be in the agenda of current models of analogical mapping.

Acknowledgments

This work was supported by Grant PICT33150 from the ANCyT of Argentina and Grant P030 from the University of Buenos Aires.

References

Falkenhainer, B. (1990) Analogical interpretation in context. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 69-76). Hillsdale, NJ: LEA

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence, 41*, 1-63.

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

Gentner, D. (1989). The mechanisms of analogical transfer. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning*. Cambridge, UK: Cambridge University Press.

Gentner, D., Holyoak, K. J., & Kokinov, B. N. (Eds.) (2001). *The analogical mind: Perspectives from cognitive science.* Cambridge, MA: MIT Press.

Gentner, D., & Kurtz, K. (2006). Relations, objects, and the composition of analogies. *Cognitive Science*, *30*, 609-642.

Hofstadter D. R., & the Fluid Analogies Research Group (1995). *Fluid concepts and creative analogies: Computer models of the fundamental mechanisms of thought.* New York: Basic Books.

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

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

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*, 427-466.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, *110*, 220-264

Kurtz, K. J. (2005). Re-representation in comparison: Building an empirical case. *Journal of Experimental & Theoretical Artificial Intelligence*, *17*, 447-459.

Reeves, M. L., & Weisberg, R. W. (1994). The role of content and abstract information in analogical transfer. *Psychological Bulletin*, *3*, 381-400.

Yan, J., Forbus, K., & Gentner, D. (2003). A theory of rerepresentation in analogical matching. *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society* (pp. 1265–1270). Mahwah, NJ: LEA