The Role of Schema-Governed Relational Categories in Analogical Inference

Ricardo A. Minervino (minervinora@gmail.com)
IPEHCS-CONICET-UNCo, Av. Argentina 1400, 8300 Neuquén
Department of Psychology, University of Comahue.

Adrián Margni (adrianmargni@gmail.com)
Department of Psychology, University of Comahue,
Irigoyen 2000, 8324 Cipolletti.

Máximo Trench (maximo.trench@crub.uncoma.edu.ar)
Department of Psychology, University of Comahue,
IPEHCS-CONICET-UNCo, Quintral 1250, 8400 Bariloche.

Abstract

The standard approach posits that analogical inferences are generated by *copying* unmapped base relations, *substituting* base entities by their corresponding target ones, and *generating* slots for unmapped base entities. Contra this account, results from Experiment 1 revealed that analogical inferences seldom include relations that resemble the base relation from which they were derived. Most of the inferences, however, could be categorized as exemplars of a schema-governed category capable of characterizing the base information to be projected. To gather further precision about the criteria that guide inference generation, in Experiment 2 we showed that analogical inferences tend to match the base information from which they are derived in values of salient dimensions of the relational category to which they belonged. Our results suggest that the relational constructs employed in modeling analogical inference should move beyond one-term multiplace predicates so as to include more complex relational structures.

Keywords: analogy; inference; relational category.

Introduction

Analogical thinking is a central mechanism in human cognition (Gentner, 2003; Hofstadter & FARG, 1995; Holyoak & Thagard, 1995), playing an important role in activities as diverse as categorization, problem solving, scientific discovery, decision making, and argumentation (Gentner, Holyoak, & Kokinov, 2001). In all these activities, analogy involves establishing a mapping between the compared situations and transferring new knowledge from a more familiar situation (*base analog*) to a less familiar one (*target analog*).

Almost all current theories of analogy agree that the alignment that takes place during mapping should satisfy the constraints of *one-to-one mapping* and *parallel connectivity,* (e.g., Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1989, Holyoak & Thagard, 1989; Hummel & Holyoak, 1997). While one-to-one mapping requires that each element in one situation maps to at most one element in the other situation, parallel connectivity entails that if two predicates are matched, their arguments must be placed in correspondence according to their roles. The following analogy illustrates these constraints:

Base analog: John loved Mary and this made John give Mary a perfume

Target analog: Peter loved Susan

While one-to-one mapping implies that pairing *John* with *Peter* should prevent pairing *John* with *Susan*, parallel connectivity dictates that if *love* has been paired with *want*, *John* must be placed in correspondence with *Peter* and *Mary* with *Susan*, as agents and patients, respectively, of the previous matched relations.

Theories of analogy also agree in that base assertions that are connected to the collection of mapped elements but which do not have a counterpart in the target will be brought over as candidate inferences. To formulate these inferences, the cognitive system would apply some variant of a *copy with substitution and generation* mechanism (CWSG; e.g., Falkenhainer et al., 1989; Holyoak, Novick, & Melz, 1994; Hummel & Holyoak, 2003). In our example, after projecting the base higher order relation *cause* to the target, the cognitive system would generate and transfer a "template" proposition like "Peter [give or a *similar action*[1]] Susan [*something like* a perfume]" into the target. This template proposition is generated from "John gave Mary a perfume", via copying *give*, substituting *John* by *Peter* and *Mary* by *Susan* (matched during mapping), and generating a slot for the entity *perfume*.

[1] The structure-mapping theory and the multiconstraint theory postulate that the new hypothetical entities should be capable of filling the role played by their corresponding base objects, but they do not give further specifications about how to identify these entities in the target domain (see, e.g., Falkenhainter et al., 1989; Holyoak et al., 1994). With respect to the transferred relations, the structure-mapping theory (see, e.g., Falkenhainer et al., 1989) maintains that a generated relation in the target will be assumed to be *identical* to the corresponding source predicate. However, to the extent that this theory has always treated identicality as a "tiered" condition (Forbus, Ferguson, Lovett, & Gentner, in press), we suppose the theory would admit inferences that include non-identical relations as long as they can be regarded as identical at a higher level of abstraction. The multiconstraint theory adheres to the "copy of relations" postulate, but it treats it as a default rule that is adequate for initial explorations in the target (see, e.g., Holyoak et al., 1994). Under these considerations, we assume that CWSG involves either copying the base relations or replacing them by similar ones.

Supporting the postulations of the dominant theories of inference generation, there is some evidence that people tend to apply the syntactic constraints of *one-to-one mapping* (i.e., previous correspondences will dictate substitutions in CWSG) and *systematicity* (i.e., people are more likely to import an inference from base to target when the fact is causally connected to other matching facts) (e.g., Clement & Gentner, 1991; Gentner, Ratterman, & Forbus, 1993; Markman, 1997). There is also evidence that people tend to apply pragmatic criteria to derive their inferences (Spellman & Holyoak, 1996). The application of these principles would be in the service of guaranteeing some minimal initial plausibility and relevance for the generated inferences (see, e.g., Holyoak et al., 1994).

As a pattern-completion process that takes maximal advantage of the mapping process, CWSG can be regarded as fast and computationally inexpensive. As most theories agree, analogical inference mechanisms should not be required to provide an adequate content to the produced inferences, something that should allegedly rely on post-inference stages of analogical reasoning such as *evaluation* and *adaptation* (Holyoak et al., 1994). The question arises as to whether the information conveyed by the templates obtained via CWSG can adequately guide post-inference generation processes in filling them in a semantically appropriate way.

In our example, if we repeat *give*, there is some probability of generating semantically appropriate inferences from the template "Peter gave [*something like* a perfume] to Susan" simply by replacing *perfume* by another exemplar of toiletry. However, while substituting *perfume* with *spa set* gives rise to a somewhat adequate inference, substituting *perfume* with *deodorant* would be inappropriate. It seems that combining *give* with an exemplar of toiletry will be adequate only if it gives rise to an instance of say, "manifestations of love". In this sense, the strategy of repeating the base relation and searching for a new exemplar of the base entity categories seems insufficient to guarantee some minimal semantic appropriateness of analogical inferences, to the extent that it requires some kind of "semantic supervision" from more complex category structures that the analogizer should keep in mind during the process. Combining substitutes (similar relations) of the relation *give* with substitutes of the entity *perfume* so as to obtain cases of "manifestations of love" would require even more thoughtful control. Just to exemplify, if we replace *give* by *lend,* no toiletry seems appropriate to generate a demonstration of love, although we can generate an exemplar of this category via replacing *perfume* by "his new car" or by "his weekend house at the beach".

The second problem with the CWSG strategy is that it can lead to inconvenient fixations, since many combinations of non similar verbs and objects could result in manifestations of love: write her a poem, prepare her favorite meal, or pick some wildflowers. Searching for cases of manifestations of love without the semantic restrictions imposed by the mechanism of CWSG seems to be a more flexible and productive strategy.

Based on the above considerations, we propose an alternative to CWSG which consists in categorizing the base analog information from which the inference will be derived as an exemplar of a schema-governed relational category (SGC), and searching for new exemplars for this category. Members of SGCs such as *murder* share a structure that can be instantiated by very different exemplars (Gentner & Kurtz, 2005; Goldwater, Markman, & Stilwell, 2011; Markman & Stilwell, 2001), such as "Fred thrust a knife into Gina's heart", "Mary had Bob drink poison", or "The offender disconnected the patient's oxygen supply". When the situations in an analogical comparison are exemplars of a SGC, the similarity between the relations and entities of the compared events is no longer necessary to have a good analogy (Minervino, Oberholzer, & Trench, 2013). The analogical relatedness between "John gave Mary a perfume" and "Peter wrote Susan a poem" is not based on semantic resemblances between *give* and *write* or *perfume* and *poem,* but rather on the fact that both acts represent exemplars of the SGC "manifestation of love". In this sense, the limitation of CWSG seems to stem from treating analogical inference as an element-by-element pattern replacement guided by isolated similarities, and from not considering the broader meaning of the facts described by propositions. When this broader meaning is taken into account, the cognitive system can do away with element-to-element similarities.

With the aim of determining which of these alternative mechanisms constitutes a better account of how analogical inferences are generated, one of the conditions of Experiment 1 served to document the extent to which analogical inferences produced by participants involve relations that are similar to those of the base analog (as posited by dominant theories), as well as the extent to which they involve facts that pertain to the same schema-governed category as the base effect. In order to confirm that participants' inferences took into account the analogical relation between the target and the source—as opposed to representing plausible consequences of the target analog considered in isolation—, the inferences generated by the abovementioned group were contrasted to the inferences produced by a second group of participants who had to propose likely consequences of the target situations, but without having previously received an analogous source.

Experiment 1

Method

Participants and Design Fifty students of psychology at the University of Comahue (mean age = 22.86 years, *SD* = 3.42 volunteered to participate in the experiment. They were randomly assigned in equal number to the analogy and the target-only groups. The dependent variables were (1) the similarity between the relation of the base effect and that of the inferred situation, and (2) whether or not the inferred situation belonged to the same SCG as the base effect on which it was inspired.

Materials Ten sets of stimuli were built, each one comprising a base and a target analog. The base analog consisted of a base cause that engenders a base effect. The base causes consisted in three-place predicates in which an agent exerts an action over an object, and directed to a patient. The base effects were predicates in which the former patient exerted another action to another object, but which is directed to the former agent. Participants were tasked with generating a consequence of the target cause that they deemed analogous to that of the base analog. Table 1 displays a sample of the experimental materials.

Table 1: Sample of experimental materials, Experiment 1

Set #	Category	Situation
1	Danger	BC: An old man left the door of his kitchen open to his two-year old grandson
		BE: The grandson ingested the old man's medicines
		TC: Another old man left the door of his kitchen open to his two-year old grandson
4	Public welfare	BC: Latvian's low-income population held a manifestation against the government
		BE: The government sent food to the low-income population b
		TC: Another low-income population held a manifestation against the government
8	Promotion of critical thinking	BC: A student questioned the theory to his professor
		BE: The professor raised the student's grade
		TC: Another student questioned the theory to his professor

Note. BC: Base Cause; BE: Base Effect; TC: Target Cause.

Procedure Participants in the analogy condition received a brief written explanation about the potential of analogical comparisons to infer new information about a target situation. The instructions presented the main activity as one in which they were going to receive a first situation comprising a cause and its associated effect, followed by a second situation for which they had to proposed an effect that could be considered analogous to that of the original fact. Participants of the target-only group received a brief written explanation about how people hypothesize effects for certain facts. The instructions presented the main activity as one in which participants were going to receive a simple situation, with the task of proposing a likely effect. Participants received the stimuli in random order. The experimental stimuli were presented on a computer screen, with participants typing their answers within prespecified fields. The administration took place in groups ranging from two to five participants, with each participant working individually. Participants were allotted a maximum of 30 min to complete the trials at their own pace.

Coding Each of the inferences proposed by participants was analyzed along two key dimensions: (1) the extent to which the action included in an inference was semantically similar to that of the base effect of the corresponding set of materials—a central prediction of the CWSG approach—, and (2) whether o not the inferred fact and its corresponding base effect belonged to the same schema-governed category. To carry out the first analysis, two judges unfamiliar with the purpose of the study received a ten-page table in which the verbs of the critical base effects (one from each set of materials) were matched against the verbs of all the inferences generated by participants for that particular set. Judges were asked to rate the similarity of the verb-pairs using a 5-point scale (1 = highly dissimilar; 5 = highly similar). They worked independently of one other, and the scores given by the two judges to each of the verb-pairs were averaged. Judges' scores were found to be reasonably reliable, Cronbach's $\alpha = .797$. While verbs yielding an average score of three or more were classified as "similar", those obtaining an average score of less than three were sorted as "dissimilar". In order to perform the second analysis, two additional judges received each of the inferences proposed by participants preceded by its corresponding target cause and followed by a list of four words or brief descriptors representing SGCs, with the instruction to draw a mark next to any of the descriptions that could be used to categorize the target effect (they could mark as many as they wanted, or leave all of them unchecked in case they considered that none of them applied). For all inferences corresponding to a given set of materials, the list of event categories comprised two SGCs that corresponded to the base effect and two SGCs that did not correspond to the base effect, all presented in random order. For example, for Set 1 (see Table 1), one of the participants generated the inference "*The grandson played with the stove*". In order to determine whether this inference could be encompassed by the same SGC as the base analog in which it was inspired, judges received the target cause plus the inference at stake, coupled with the following event descriptions: (1) revenge, (2) dangerous situation (3) jealousy reaction, and (4) risky situation. Inferences were scored as sharing a SGC with the base effect in all those cases where the two judges checked at least one of the two "correct" event descriptors (danger and/or risk), regardless of whether they agreed on which of the correct descriptors was checked.

Results and Discussion

The verbs of the inferences generated by participants of the analogy group resembled those of their corresponding base effect in 31.6% of the cases. It should be noted, however, that 12.7% of the inferences generated in response to the target analog alone (i.e., those of the target-only group) involved verbs that resembled those of the base effect received by participants of the analogy group. More fine-grained analyses using chi-square statistics revealed that the rate of utilization of similar verbs by the analogy group differed from the rate of spontaneous utilization of those same verbs by the target-only group in 4 of the 10 sets of materials (see Table 1).

With regards to SGC similarity, judges' analyses showed that while the inferences generated by the analogy group involved the SGC of the base effect in 88% of the cases, the inferences produced in response to the target analog alone belonged to these same categories in 34.4% of the cases. Chi- square tests revealed that for all 10 sets of materials the probability of generating an inference that pertains to the SGC of the base effect by participants of the analogy group was higher than the proportion of inferences pertaining to those same SGCs within the target-only condition (See Table 1).

The low proportion of semantically similar relations among the inferences produced by the analogy group suggests that the mechanism of postulating target relations that resemble their counterparts in the base analog, as dictated by CWSG, cannot adequately account for how analogical inferences are derived. In contrast, the fact that the vast majority of the inferences belonged to the same SGC as the causal consequent of the base analog suggests that the dominant mechanism involved in the generation of analogical inferences consists in analyzing the SGCs to which the base effect belongs, and generating further exemplars of such categories.

Table 2. Percentages of inferences exhibiting verb and relational category similarity with the base analog

	Similar relations			Same relational category		
Set #	Analogy condition	Target-only condition	χ^2	Analogy condition	Target-only condition	χ^2
1	32% (8)	16% (4)	1.75	92% (23)	36% (9)	17.01**
2	36% (9)	12% (3)	3.95*	88% (22)	40% (10)	12.5**
3	32% (8)	8% (2)	4.5*	84% (21)	52% (13)	5.88*
4	28% (7)	12% (3)	2	88% (22)	12% (3)	28.88**
5	20% (5)	12% (3)	0.6	84% (21)	24% (6)	18.12**
6	32% (8)	36% (9)	0.09	92% (23)	20% (5)	26.3**
7	16% (4)	36% (9)	2.6	88% (22)	56% (14)	6.35*
8	44% (11)	0% (0)	14.1*	88% (22)	36% (9)	14.35**
9	44% (11)	4% (1)	10.96*	96% (24)	32% (8)	22.22**
10	32% (8)	36% (9)	0.09	80% (20)	28% (7)	13.61**

Note. *Significant at $\alpha = .05$; **Significant at $\alpha = .01$

Having documented that the majority of inferences were exemplars of a relational category that was readily applicable to the base analog effect, a sensible research question concerned whether inclusion to such relational category suffices as a criterion for generating analogical inferences. As suggested by data obtained by Minervino et al. (2013), a factor that seems to influence the perception of analogical resemblance between exemplars of a SGC has to do with whether the target situation matches the base situation along the most salient dimensions of the relational category to which they belong. Taking the category *robbery* as an example, the analogability of two exemplars depends on whether they match in central dimensions such its importance, violence or planning. Experiment 2 was aimed at determining whether the observed sensitivity to this constraint generalizes to analogical inference.

As in Experiment 1, participants received a base analog comprising two causally related situations, followed by the presentation of a target situation that was virtually identical to the causal consequent of the base analog and by the task of completing the target situation with a consequence that they deemed analogous to that of the base. The main difference with Experiment 1, however, was that the exemplars of SGCs that were employed as the effects of the base situations were chosen to instantiate either a high or a low value along a central dimension of the relational category to which they belonged. The purpose of the experiment was to assess the extent to which the exemplars of SGCs included in participants' inferences matched the consequent of the base situation in terms of its values along the manipulated dimension.

Experiment 2

Method

Participants Twenty-four students of psychology at the University of Comahue (mean age = 21.1 years, *SD* = 3.36) volunteered to participate in the study.

Materials and Procedure Ten new sets of materials were built. The sets had the same general structure as those of Experiment 1, with the main difference being that for each base cause we derived two possible base consequences instead of one. These two consequences belonged to the same SGC, but differed from each other in that they scored differently along a central dimension of such category. As an example, the base cause *A paleontologist brought fossils to the Trelew Museum* was followed either by the base effect *The museum commissioned a statue of the paleontologist* (an instance of *reward* with a high value in the dimension "magnitude") or by *The museum issued a diploma to the paleontologist* (low value in the dimension "magnitude"). Table 3 displays a sample of the experimental materials.

To ensure that participants encoded the base consequences as members of the SGC whose critical dimension was being manipulated, participants were explicitly informed about the specific category to which the base consequence belonged (see Table 3).

Table 3: Sample of experimental materials, Experiment 2

Set	Category	Base and target situations
1	REWARD	BC: A paleontologist brought fossils to the Trelew Museum BE_{hv}: The Museum commissioned a statue of the paleontologist BE_{lv}: The Museum issued a diploma to the paleontologist TC: Another paleontologist brought important fossils to the Rawson Museum
2	ROBBERY	BC: The old lady trusted her house's keys to her nanny BE_{hv}: The nanny sold the old lady's jewelry BE_{lv}: The nanny took a book from the old lady's house TC: Another old lady trusted her house's keys to her nanny
5	CONTRI-BUTION	BC: A young man was invited to a barbecue by his friends BE_{hv}: He volunteered to pay the meat to his friends BE_{lv}: He volunteered to bring matches to his friends TC: Another young man was invited to a barbecue by his friends

Note. BC: Base Cause; BE_{hv}: Base effect with high values on a key dimension of the relational category; BE_{lv}: Base effect with low values on such dimension; TC: Target Cause.

Two complementary booklets of materials were built. In each version half of the sets were coupled with consequences instantiating low values along the manipulated dimension of the SGC to which they belonged, and half with consequences embodying high values along such dimensions. The procedure was identical to that on Experiment 1.

Coding Two new judges received each of the inferences generated by participants coupled with the critical dimension that corresponded to that set of materials. They were asked to rate how it fared along such dimension using a 5-point scale ranging from the minimum to the maximum possible levels along the manipulated dimension (e.g., for the reward example, they had to rate the *magnitude* of the reward from 1 = very small, to 5 = huge). The scores given by the two judges to each of the inferences were averaged. Judges' scores were found to be reliable, Cronbach' $\alpha = .823$.

Results and Discussion

The inferences generated out of base facts ranking high along the manipulated dimensions obtained higher scores than those generated out of base facts displaying lower levels along that dimension ($M = 3.18$, $SD = 0.40$ vs. $M = 2.18$, $SD = 0.36$, $t(25) = -10.05$, $p < .001$.

These results demonstrate that the way in which the base effect fares along a critical dimension of the SGC to which it belongs constrains the way analogical inferences will fare along such dimension. In order to gather a subtler estimate of the strength of this association, judges were also required to score the base effects along the manipulated dimensions. The correlation between the scores assigned to the base effects and those of their associated inferences was strong, $r = .476$, $n = 260$, $p < .001$. Furthermore, in 71.54% of the cases the scores of the generated inferences along the manipulated dimension were no farther than one point away from those of the base analogs on which they were inspired.

General Discussion

A key prediction of the CWSG approach to analogical inference consists in that people will tend to construct their inferences repeating the base relations from which the inferences will be derived or replacing them by similar ones. Against this position, Experiment 1 showed that people do not care much about preserving similarity to base relations but instead focus on generating new exemplars of the SGC applied to the base information from which inferences are derived. In Experiment 2 we collected data favoring a further hypothesis associated to our category-based perspective, namely, that when proposing new exemplars of SGCs people tend to generate cases that fare closer to the base exemplar along critical dimensions of the category to which they pertain.

We have argued that the templates generated by CWSG could sometimes be insufficient to guide the analogizer in generating semantically sensible inferences during the post inference stages of evaluation and adaptation. Returning to the example presented in the Introduction, the chances of generating semantically appropriate inference from the template "Peter [give or a *similar action*] Susan [*something like* a perfume] seems rather low. We speculated that while some toiletries could perhaps result in a sensible inference, others do not, and that post-inference generation mechanisms have no semantic basis to distinguish between them.

The standard approach to analogical inference generation could argue that *"something* like a perfume*"* should not be interpreted as "an exemplar of toiletry", and that this interpretation is to some extent caricaturizing CWSG, since an intelligent system operating in an analogical mode will not be guided by superficial similarities such as membership to a same category, but would rather interpret it as, say, "give + things that a woman finds romantic". In this sense, the system would promote the search for new exemplars of this *ad hoc* category (e.g., a teddy, necklace or a bouquet). The problem with this argument is that the very consideration of this *ad hoc* category supposes the prior conceptualization of the template as a "manifestation of love", something that the analogical machinery has not generated. It is possible that the generalized support that the CWSG approach has received comes in part from the fact that programmers inadvertently read far more understanding than is warranted into the templates produced by this mechanism, as an effect of projecting the SGC that they apply to capture the whole meaning of the template.

The standard approach could also argue that the analogical machinery was not meant to deal with the activity of comprehending the analogs, but rather to start operating once the analogs have been fully comprehended (see, e.g., Morrison & Dietrich, 1995). In this vein, the analogical engine would receive the fact that *John gave Mary a perfume* already interpreted as a case of "carrying out a very romantic manifestation of love" (i.e., as a case of this category, and with some specific properties). A problem with this argument is that this conceptualization cannot be captured by a relation—defined as *one-term* multiplace predicates. The execution of a very romantic demonstration of love would be propositionally represented, *stricto sensu*, as CARRY OUT [John, ((VERY) ROMANTIC (manifestation of love)), Mary]. The essential information to be transferred is located in an argument represented as a noun (*manifestation of love*) and its property (VERY ROMANTIC), and not in the one-term predicate outside the brackets (CARRY OUT).

We are far from calling into question the importance of relational aspects in analogical thinking, but we believe it is necessary to discuss and amplify the meaning of "relational", so as to avoid reducing it to one-term multiplace predicates. It should be broadened to include, for example, *relational structures* as those captured by SGCs. In these structures, relations are only a constituent, being other thematic roles (e.g., agents, patients, objects or instruments) just as important. For example, if an instance of the category "manifestation of love" includes the relation *give*, the agent's intention should be to awake certain emotions in a person, the patient has to be a candidate for being emotionally affected by the agent at stake, and the object should be pleasant to the patient. The complex interdependency of the constituents of a fact that make it pertain to a SGC makes it proper to talk about these categories as "relational" structures, but the sense of the term is broader than the one employed in computational models of analogy (i.e., a one-term multiplace predicate). The relational character of these categories is also evident in the extremely different situations that can constitute exemplars of a SGC, which could differ even in their relations defined in the traditional way.

Acknowledgments

We thank Micah Goldwater for suggesting how to operationalize whether participants' inferences belonged to the SCG of the source. This work was supported by the National Agency for Scientific and Technical Research (ANPCyT) under Grants PICT 2352 and PICT 2650, by the University of Comahue under Grants B213, C123 and C134, and by the National Council for Scientific and Technical Research (CONICET) under Grant PIP 0567.

References

Clement, C. A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science*, *15*, 89-132.

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

Forbus, K. D., Ferguson, R. W., Lovett, A., & Gentner, D. (in press). Extending SME to handle large-scale cognitive modeling. *Cognitive Science.*

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. (2003). Why we are so smart. In D. Gentner, & S. Goldin-Meadow (Eds.)*, Language in mind: Advances in the study of language and thought*. Cambridge, MA: MIT Press.

Gentner, D., & Kurtz, K. J. (2005). Learning and using relational categories. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. W. Wolff (Eds.), *Categorization Inside and Outside the Laboratory* (Vol. 43). Washington, DC: APA.

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

Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, *25*, 431-467.

Goldwater, M. B., Markman, A. B, & Stilwell, C. H. (2011). The empirical case for role-governed categories. *Cognition*, *118*, 359-376.

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

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.

Markman, A. B. (1997). Constraints on analogical inference. *Cognitive Science*, *21*, 373-418.

Markman, A., & Stilwell, C. (2001). Role-governed categories. *Journal of Experimental & Theoretical Artificial Intelligence*, *4*, 329–358.

Minervino, R., Oberholzer, N., & Trench, M. (2013). Overall similarity overrides element similarity when evaluating the quality of analogies. *Journal of Cognitive Science*, *14*, 287-317.

Morrison, C., & Dietrich, E. (1995). Structure-mapping vs. High-level perception: The mistaken fight over the explanation of analogy. In J. D. Moore and J. Fain Lehman (Eds.), *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 678-682). Hillsdale, NJ: Erlbaum.