

# Absence Makes the Thought Grow Stronger: Reducing Structural Overlap Can Increase Inductive Strength

Hee Seung Lee (heeseung@ucla.edu)

Department of Psychology, University of California, Los Angeles  
Los Angeles, CA 90095 USA

Keith J. Holyoak (holyoak@lifesci.ucla.edu)

Department of Psychology, University of California, Los Angeles  
Los Angeles, CA 90095 USA

## Abstract

Computational models of analogy have assumed that the strength of an inductive inference about the target is based directly on similarity of the analogs, and in particular on shared higher-order relations. However, in Experiment 1 we show that reducing analogical overlap by eliminating a higher-order causal relation (a preventive cause present in the source) from the target *increased* inductive strength even though it *decreased* similarity of the analogs. In Experiment 2 we extend these findings to cross-domain analogical inferences based on correspondences between higher-order causal relations. Analogical inference appears to be mediated by building and then “running” a causal model. We discuss the implications of the present findings for theories of both analogy and causal inference.

**Keywords:** analogical inference; causal model; similarity; inductive strength; preventive cause; generative cause.

## Introduction

When people face uncertainty, they often make analogical inferences about an unfamiliar *target* situation based on a similar but better understood *source* situation. However, people do not draw inferences from the source about all the missing target properties. For example, consider the earth as a source analog for inferring properties of the moon. It seems more likely the resulting analogy will lead to the inference that the moon may contain iron deposits than that the moon has a system of freeways, even though the earth contains iron deposits and also has freeways. In general, inductive inferences seem to be guided by certain general constraints that allow people to make analogical inferences *selectively* (Holyoak, 1985).

### Inductive constraints: Are causal relations special?

People are more likely to draw confident analogical inferences when they perceive a source and target to be similar, particularly if they overlap in their relational structure. According to Gentner’s (1983) *systematicity* principle, higher-order relations are especially important in promoting analogical inferences. The canonical example of a higher-order relation is “cause”; however, the “order” of a relation is defined syntactically (a relation that takes another relation as an argument), rather than by meaning.

A central issue in understanding constraints on analogical inference is whether causal relations have some special

status by virtue of the fact that causes actually produce (or sometimes prevent) their effects (e.g., Bartha, in press; Hesse, 1966; Winston, 1980), or whether causal relations simply operate as special cases of higher-order relations, defined by the formal properties of predicate-argument structures. Holyoak (1985) emphasized that causal knowledge is the basis for *pragmatic* rather than just structural constraints on analogical inference. Using complex stories, Spellman and Holyoak (1996) showed that when the source-target mapping was ambiguous by structural criteria, those relations causally-relevant to the reasoner’s goal determined the preferred mapping and inferences about the target. Lassaline (1996) demonstrated that people make stronger inferences on the basis of the higher-order relation “cause” than on the basis of a non-causal relation, “temporally prior to”, which appears to have the same formal structure. Thus the same syntactic “order” of relations does not always yield the same degree of inductive strength about the target property to be inferred.

There is no doubt that causal relations play an important role in analogical inference. However, very little effort has been devoted to incorporating basic elements of *causal models* (Griffiths & Tenenbaum, 2005; Waldmann & Holyoak, 1992), such as the distinction between generative and preventive causes (Cheng, 1997; Lu et al., 2007), into models of analogy. Recently, Lee and Holyoak (2007) examined the role of causal models in analogical inference by testing a hypothesis proposed by Bartha (in press). Bartha distinguished between *contributing* causes (generative) and *counteracting* causes (preventive) in evaluating the strength of analogical arguments. He pointed out that eliminating from the target a counteracting cause present in the source might actually strengthen an argument from analogy. For example, iron deposits are still present on earth despite the fact that humans have been extracting iron ore in mining operations for centuries. Taking the earth as a source analog for the moon, the fact that no mining operations have so far been conducted on the moon (a mismatch with a property of the source) seems to *strengthen* the analogical inference that iron deposits remain to be found on the moon.

To test this hypothesis, Lee and Holyoak examined use of causal models by adapting a paradigm introduced by Lassaline (1996). Participants read a description of two

imaginary animals simply referred to as *Animal A* and *Animal B*, and then evaluated either strength of an analogical inference or similarity of the two animals. Across arguments, presence versus absence of a generative or else preventive relation connecting one shared property to a non-shared property was manipulated. Lee and Holyoak found that similarity ratings increased with the inclusion of a shared property; but whereas inductive strength judgments increased with the inclusion of a shared generative property, judged strength increased with the *removal* of a shared preventive property. These results suggest that people evaluate whether the causal relations in the source are generative or preventive, and then use the resulting causal models to infer the probability of obtaining a corresponding effect in the target.

However, it could be argued that the particular causal relations used in Lee and Holyoak's (2007) study held between simple attributes of a single entity, rather than propositions, and hence were actually first-order relations as defined by Gentner (1983). Also, the materials used in the experiments all involved within-domain analogical transfer from one animal to another animal of the same species. In contrast, the core examples of analogical reasoning discussed in the literature involve long-distance, cross-domain transfer between situations based on highly dissimilar entities that nonetheless share relations. In the present study, we examined whether dropping a shared preventive cause would increase rather than decrease the strength of an analogical inference when the cause was clearly a high-order relation (Experiments 1-2) and when the analogy required cross-domain transfer (Experiment 2).

## Experiment 1

Experiment 1 was designed to investigate whether or not people use causal models during analogical inference when the source includes more complex systems of higher-order relations. Adapting the paradigm from Lee and Holyoak (2007), we developed a set of materials in which cause-effect relations held between relations connecting multiple objects. Schematically, the form of the causal relations we used was: [A is greater than B] causes [E]. In this example, there are three objects A, B, and E. The objects A and B enter into a first-order relation (i.e., *greater than*). This first-order relation between the objects A and B is causally connected to effect E. The structure of such a causal relation clearly meets Gentner's (1983) definition of a higher-order relation, as it takes a proposition (e.g., *A is greater than B*) as an argument. By investigating whether people use causal models in evaluating analogical inferences based on this more complex type of structure, it will be possible to verify the general role of causal models as a major inductive constraint on analogical inference.

## Method

**Participants** Sixty undergraduate students at the University of California, Los Angeles (UCLA) participated in the experiment for course credit. Half of the participants

provided inductive strength judgments, and the other half provided similarity ratings.

**Materials and Procedures** Participants read a description of a fanciful "newly discovered species of bird", which specified three causal relations between the relative amount of hormone, enzyme, and neurotransmitter (causes) and an abnormal characteristic (effect) found in the bird. The relative amount of two substances of each of two types was described as tending to produce the abnormal characteristic, and the relative amount of two substances of the third type was described as tending to prevent the abnormal characteristic. An example of descriptions is the following:

[Hormone A > Hormone B]  
tends to PRODUCE blocked oil glands.  
[Neurotransmitter X > Neurotransmitter Z]  
tends to PRODUCE blocked oil glands.  
[Enzyme P > Enzyme Q]  
tends to PREVENT blocked oil glands.

After reading the description of findings for the first described bird, participants studied reports about the relative amounts of hormones, neurotransmitters, and enzymes found in a second bird of the same species. Four different types of argument were created:  $G_1G_2X$ ,  $G_1G_2P$ ,  $G_1G_2$ , and  $G_1P$ . G and P represent generative and preventive causes, respectively. The factor X was constructed by switching the order of objects in the P relation; e.g., in the above example, X was [Enzyme Q > Enzyme P].

Eight different descriptions of birds were created, and each argument type was assigned to each description, creating 32 items altogether. Of the total of 32 items available, eight items were used to create a booklet for each participant, two of each argument type ( $G_1G_2X$ ,  $G_1G_2P$ ,  $G_1G_2$ , and  $G_1P$ ). This counterbalancing generated four different sets of materials, thereby avoiding repeated use of the same abnormal characteristic for an individual participant. Within each set, the order of items was randomized for each participant.

For the group making similarity judgments, the task was to rate how similar the two birds are, choosing a number between 0 (totally different) and 10 (identical). For the group making inductive strength judgments, the task was to judge how likely a second bird would have the abnormal characteristics described in the first bird, based on the descriptions of the relative amounts of hormones, neurotransmitters, and enzymes found in the two birds. In making their judgments, participants were asked to imagine there were 100 examples of the second bird, and to estimate how many out of these 100 cases would have the property stated as the conclusion, assigning a number between 0 and 100 for each item.

## Results and Discussion

The results for both similarity ratings (left) and inductive strength judgments (right) are shown in Figure 1. Similarity ratings and inductive strength judgments were analyzed

separately. For each dependent measure, a one-way ANOVA was performed with the four argument types ( $G_1G_2X$ ,  $G_1G_2P$ ,  $G_1G_2$ , and  $G_1P$ ) as a within-subjects variable.

The mean similarity ratings of analogies based on  $G_1G_2X$ ,  $G_1G_2P$ ,  $G_1G_2$ , and  $G_1P$  arguments were 6.35, 8.73, 7.57, and 5.03, respectively. These means were significantly different,  $F(3, 87) = 24.55$ ,  $MSE = 3.09$ ,  $p < .001$ .  $G_1G_2P$  arguments were rated as having the highest similarity of the four argument types.  $G_1G_2P$  arguments were rated as having higher similarity than those arguments with fewer shared relations, either those excluding P (mean of  $G_1G_2X$  and  $G_1G_2$  arguments),  $t(29) = 3.91$ ,  $p = .001$ , or those excluding G ( $G_1P$  arguments),  $t(29) = 8.50$ ,  $p < .001$ . In addition,  $G_1G_2$  was rated as having higher similarity than  $G_1G_2X$  (where X was formed by reversing the relation P in the source),  $t(29) = 3.84$ ,  $p = .001$ .

The mean inductive strength judgments for  $G_1G_2X$ ,  $G_1G_2P$ ,  $G_1G_2$ , and  $G_1P$  arguments were 87.40, 69.83, 87.47, and 43.70, respectively. The mean inductive strength judgments differed across argument types,  $F(3, 87) = 75.17$ ,  $MSE = 170.53$ ,  $p < .001$ , but showed a different pattern than that obtained for similarity ratings.  $G_1G_2P$  was evaluated as having lower inductive strength than the arguments that dropped preventive relations (mean of  $G_1G_2X$  and  $G_1G_2$  arguments),  $t(29) = 5.78$ ,  $p < .001$ , but as having higher inductive strength than the arguments that dropped generative relations ( $G_1P$  arguments),  $t(29) = 8.12$ ,  $p < .001$ . Unlike the pattern for similarity ratings, inductive strength judgments did not differ for arguments  $G_1G_2X$  and  $G_1G_2$ ,  $t(29) = .02$ ,  $p = .98$ .

To summarize, the results of Experiment 1 extended the basic findings of Lee and Holyoak (2007) to more complex causal structures in which the causal relations were higher-order, as defined by Gentner (1983). Similarity ratings increased with the addition of shared relations, regardless of whether the relation was generative or preventive. Also, the argument  $G_1G_2X$  was rated as having lower similarity than the argument  $G_1G_2$ . Even though these arguments had the same number of shared relations (two:  $G_1$  and  $G_2$ ), X was not present in the source, so this relation constituted a difference between the source and target, reducing the perceived similarity (Tversky, 1977).

The results for inductive strength judgments showed a different pattern than that observed for similarity ratings. The  $G_1G_2P$  arguments had higher inductive strength than the  $G_1P$  arguments (the same pattern as for similarity ratings), but had lower inductive strength than the arguments  $G_1G_2X$  and  $G_1G_2$  (a pattern opposite to that obtained for similarity ratings). Thus whereas the presence of a shared generative relation increases inductive strength, the *absence* of a preventive relation increases inductive strength.

In addition, unlike the results for similarity ratings, the arguments  $G_1G_2X$  and  $G_1G_2$  were evaluated as having equally high inductive strength. In the inductive strength task X was not related to the effect property, so that when inferring the presence of the effect in the target, X might

have been interpreted as a neutral property. Accordingly, the arguments  $G_1G_2X$  and  $G_1G_2$  were viewed as equally strong because they involved the same shared causal relations,  $G_1$  and  $G_2$ .

The findings from Experiment 1 confirm that people do not simply focus on the number of correspondences between a source and target; rather, they consider the *meaning* of causal relations, and in particular whether causes are generative or preventive. By using materials involving causal relations at a higher level of formal complexity, the results of Experiment 1 show that people use causal models during analogical inference even when the source involves complex higher-order relations among multiple objects.

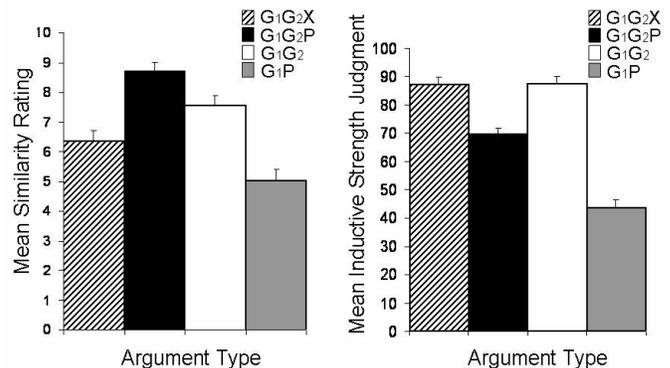


Figure 1: Mean similarity ratings (left) and mean inductive strength judgments (right) for each argument type in Experiment 1.

## Experiment 2

In Experiment 1, we manipulated features of animals in a source and target, and assumed that participants used analogical reasoning to infer a certain missing property in the target based on the source information. However, it is possible that when reasoning about new instances of the well-known category of animals, people actually reasoned at the level of categories, essentially assuming that all individuals of the same animal species share the same basic biological characteristics. Given information about the source animal, people may have naturally projected its properties to the category of the novel species, and then in turn applied this categorical knowledge to make inferences about the target animal of the same species. That is, given a source animal with the structure  $G_1G_2P$ , people may infer that these causal relations hold for all animals of that species, in effect setting up a causal model for the species as a whole, rather than directly modeling the causal structure of the target individual based on the source. (See Rehder, 2007, for a review of work relating causal models to category-based induction.)

Accordingly, Experiment 2 was performed to examine whether dropping a higher-order relation (a preventive cause) present in the source can increase inductive strength when the source and target are drawn from distinct semantic domains.

## Method

**Participants** Fifty-two undergraduate UCLA students received course credit for participating in the experiment. Participants were randomly assigned to one of two conditions,  $G_1G_2P$  or  $G_1G_2$ .

**Materials and Procedures** Two stories were created, one serving as a source story and the other as a target story, based on two different domains, chemistry and astronomy. A chemist's observations about three liquids and an astronomer's observations about three stars served as the source and target story, respectively. To ensure that people could not use their prior knowledge of liquids or stars to make inferences, all the liquids and stars were novel and imaginary.

Participants first read the source story:

A research chemist has recently discovered that if three liquid substances, Denitrogeel, Oreor, and Tetosium, are mixed together, then a chemical change sometimes occurs so that the molecules of Denitrogeel and Tetosium bond together. With molecules of Oreor serving as a catalyst, the molecules of Denitrogeel and Tetosium attract each other and the mixed liquid becomes very adhesive, finally changing into a solid material. Through repeated experiments, the scientist has also identified three main factors that determine whether or not the mixed liquids change into a solid.

Following this cover story, three relational observations about the liquids were listed, as summarized in Table 1. Two of the relations tended to cause the effect, changing mixed liquids into a solid, and one of the observations tended to prevent this effect. For example, the fact that "Denitrogeel is colder than Oreor" tended to cause formation of a solid, whereas the fact that "The volume of Tetosium is greater than the volume of Denitrogeel" tended to prevent formation of a solid. The causal relations used in Experiment 2, like those in Experiment 1, were formally higher-order relations. After reading the source story, participants read an astronomer's observations about three stars:

An astronomer who reads about the chemist's findings thinks it may be possible that three stars, Acruxia, Errailel, and Castoriff, located in a distant galaxy, behave in a way similar to the three liquids. The theory is that gravitational attraction among all three stars could make two of the stars move closer together, so that two stars finally fuse to form a super-star. The

three stars are close to each other and no other stars have been found in that region of the galaxy.

Following this description of the astronomer's hypothesis, two or three facts about these stars were listed (see Table 1). All of these facts were structurally parallel with the chemist's findings. Table 1 shows the correspondences between the two stories. The corresponding relations across the two stories had varying degrees of semantic overlap. The semantic similarities between relations in the two stories (e.g., "being more turbulent" and "being subject to more violent solar storms") were intended to facilitate the process of mapping the two analogs and identifying their relational correspondences.

The experimental design included just two conditions,  $G_1G_2P$  and  $G_1G_2$ . As in the previous experiment, G and P represent a generative cause and a preventive cause, respectively. In the  $G_1G_2P$  condition, two of the facts in the target were semantically and structurally consistent with the corresponding causes of the effect in the source, and one of the target facts was semantically and structurally consistent with the preventive relation in the source. In the  $G_1G_2$  condition, there was no match in the target to the preventive cause in the source. To prevent a blind mapping between the source and target based on the order of listed facts, the order of the facts was randomized for each participant.

After reading the two stories, the chemist's liquid story and the astronomer's star story, two tasks were administered to all participants. We did not collect ratings of similarity in Experiment 2, as it is transparent that the analogs are more similar in the  $G_1G_2P$  (three shared relations) than in the  $G_1G_2$  condition (two shared relations). We did, however, assess participants' ability to identify the relational correspondences between the two stories. In this mapping task (always administered first), participants were asked to identify which of the three stars (Acruxia, Errailel or Castoriff) corresponded most closely to each of the three liquids (Denitrogeel, Oreor, and Tetosium). This mapping task was intended to ensure that participants read the two stories, and to assess whether they could in fact identify structural parallels between the source and the target (a prerequisite for drawing sensible analogical inferences).

The second task involved drawing an analogical inference about the astronomy problem based on the source analog from the domain of chemistry. Participants were asked to judge how likely it was that two of the three stars would fuse to form a super-star. To answer this question, they were

Table 1: Schematic structure of analogies used in Experiment 2.

Story	Source: Chemist's observations about three liquids	Target: Astronomer's observations about three stars
$G_1$	Denitrogeel is colder than Oreor	Acruxia has lower temperature than Errailel
$G_2$	Oreor is stirred vigorously so it is more turbulent than Tetosium	Errailel is subject to more violent solar storms than is Castoriff
P	The volume of Tetosium is greater than the volume of Denitrogeel	The diameter of Castoriff is wider than the diameter of Acruxia
Effect	Changing into a solid	Forming a super-star

told to assume that everything in the descriptions was true and to focus on analogous relations between the chemist's observations and the astronomer's observations. This inductive judgment was made using the same type of frequency scale as used in Experiment 1. Participants were told to imagine that the astronomer observes 100 cases, and to estimate in how many of these cases two of the three stars would fuse to form a super-star, giving a number between 0 and 100.

## Results and Discussion

If participants failed to correctly answer all three questions about liquid/star correspondences, we coded their mapping performance as incorrect. Six out of the 52 participants (four in the  $G_1G_2P$  condition and two in the  $G_1G_2$  condition) gave incorrect mappings. Because incorrect mappings could have led to erroneous analogical inferences, we first report analyses of inference data excluding these participants. Mean inference judgments for  $G_1G_2$  and  $G_1G_2P$  arguments (excluding data from participants who made incorrect mappings) were 69.20 and 52.04, respectively (see Figure 2). These data were analyzed using an independent-samples  $t$  test, which revealed a significant effect of argument type,  $t(44) = 2.59, p = .013$ . Even when inference data from participants who gave incorrect mappings were included, the effect of argument type remained reliable,  $t(50) = 2.20, p = .032$ .

The results of Experiment 2 thus confirm and extend the basic finding obtained in Experiment 1: When a preventive causal relation present in the source analog is absent from the target, the strength of an inductive inference is correspondingly increased. By demonstrating this effect with cross-domain analogical transfer, the results of Experiment 2 make it clear that this beneficial effect of *reducing* overlap of a higher-order relation is not solely attributable to category-based induction; rather, the phenomenon is observed when people must base an inference about a target analog directly on a single source analog.

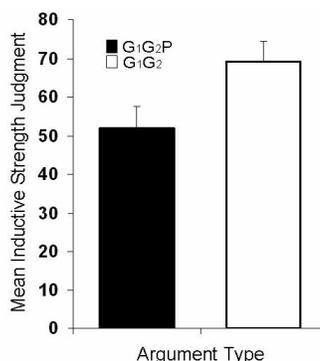


Figure 2: Mean inductive strength judgments for each argument type in Experiment 2.

## General Discussion

In Experiment 1 we showed that reducing analogical overlap by eliminating a higher-order causal relation (a

preventive cause present in the source) from the target *increased* inductive strength even though it *decreased* similarity of the analogs. In Experiment 2 we extended these findings to cross-domain analogical inferences based on correspondences between higher-order causal relations. In both experiments, the preventive causal relation was clearly a higher-order relation in the formal sense defined by Gentner (1983). Also, because in Experiment 2 cross-domain analogical transfer cannot be accounted for by inferences based on pre-existing categories (i.e., category-based induction; see Rehder, 2007), our overall findings support the conclusion that analogical inference involves using the source analog to guide construction of a causal model of the target analog.

Even though most of the major theoretical discussions of analogy, beginning with Winston (1980), have in one way or another acknowledged the critical importance of causal knowledge as a basis for analogical inference, extant computational models are not able to explain the present findings. The systematicity principle and its computational implementation, SME (Falkenhainer et al., 1989), base analogical inference solely on the logical form of representations and not on their meaning, and hence are committed to the prediction that shared higher-order relations can only help and never hinder analogical inferences. The fact that a shared higher-order relation—a preventive cause—in fact *reduces* analogical transfer thus provides a compelling demonstration that any successful model of analogy will need to deal with the meaning of semantic representations, not just their logical form.

Even computational models that allow for the special importance of causal relations are in their current implementations unable to account for our present findings. For example, both ACME (Holyoak & Thagard, 1989) and LISA (Hummel & Holyoak, 2003) include mechanisms for placing greater weight on causal relations so that these have especially strong influences on analogical mapping and inference. In their current implementations, ACME and LISA would naturally mark both generative and preventive causes as “important” and focus attention on them. However, these models are unable to grasp that the two types of causes are important in different ways, with very different implications for analogical inference—a preventive cause is indeed important, but the outcome would be more probable without it.

The problem for all extant models of analogy is that they lack a detailed representation of causal knowledge that could support commonsense causal reasoning (e.g., generative causes make things happen, but preventive causes stop things from happening; causes produce effects, but effects don't produce their causes). The “common sense” nature of the present findings highlights the weakness of commonsense reasoning in current models of analogy.

More generally, the present findings add to a growing body of work suggesting that future theoretical progress in understanding analogical reasoning requires taking a

broader view of the overall process. As pointed out by Bassok, Wu and Olseth (1995), analogy models have typically focused on the mechanisms by which correspondences between predetermined representations of a source and target can be established and then exploited to generate candidate inferences about the target. Both the input representations (source and target) and the output (target inferences) are treated as static. But as Bassok et al. argued, human analogical reasoning appears to be based on highly dynamic representations. The initial representations of the source and target may undergo semantic interpretation and elaboration that affect the ease of analogical transfer.

Similarly, the present findings call attention to the dynamic nature of the *output* of analogical processing. In current analogy models the candidate inferences, like the analogs themselves, are represented in some static code. The degree of belief in an inference is solely determined by some measure of the “goodness” of the source-target mapping from which the inference was derived. What is lacking, given the findings reported here, is some mechanism for dynamically updating the degree of belief in an analogical inference based on the causal relationships within the target itself.

Moreover, our findings suggest that the missing theoretical mechanism for dynamic inference evaluation cannot be simply “outsourced” to some post-analogical module, such as “verification” based on direct knowledge acquired about the target (as discussed, for example, by Falkenhainer et al., 1989, p. 23). Consider the situation confronting the reasoner in our Experiment 2: No direct knowledge about the behavior of the imaginary star system that forms the target domain is ever provided, beyond that knowledge allowing computation of the analogical inference in the first place. That is, the reasoner has no additional basis for assessing how likely it is that two stars will fuse to form a super-star. Rather, the causal model of the target—constructed by analogical inference from the source domain of chemical reactions—itself allows dynamic evaluation of degree of belief in the candidate inference. As Bartha (in press) argues on the basis of actual examples from the history of science, the credibility of analogical inferences can be assessed in part by internal criteria. In non-experimental sciences, such as astronomy and archaeology, analogy may at times provide the most direct source of knowledge available for evaluating inferences about the target.

### Acknowledgments

Preparation of this paper was supported by ONR grant N000140810186 to KH.

### References

Bartha, P. (in press). *By parallel reasoning: The construction and evaluation of analogical arguments*. Oxford, UK: Oxford University Press.

- Bassok, M., Wu, L.-L., & Olseth, K. L. (1995). Judging a book by its cover: Interpretative effects of content on problem-solving transfer. *Memory & Cognition*, *23*, 354-367.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367-405.
- 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.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*, 334-384.
- Hesse, M. (1966). *Models and analogies in science*. Notre Dame, IN: University of Notre Dame Press.
- Holyoak, K. J. (1985). The pragmatics of analogical transfer. In G. H. Bower (Ed.), *The psychology of learning and motivation*, Vol. 19 (pp. 59-87). New York: Academic Press.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, *13*, 295-355.
- Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, *110*, 220-264.
- Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 754-770.
- Lee, H. S., & Holyoak, K. J. (2007). Causal models guide analogical inference. In D. S. McNamara & G. Trafton (Eds.), *Proceedings of the Twenty-ninth Annual Conference of the Cognitive Science Society* (pp. 1205-1210). Austin, TX: Cognitive Science Society.
- Lu, H., Yuille, A., Liljeholm, M., Cheng, P. W., & Holyoak, K. J. (2007). Bayesian models of judgments of causal strength: A comparison. In D. S. McNamara & G. Trafton (Eds.), *Proceedings of the Twenty-ninth Annual Conference of the Cognitive Science Society* (pp. 1241-1246). Austin, TX: Cognitive Science Society.
- Rehder, B. (2007). Property generalization as causal reasoning. In A. Feeney & E. Heith (Eds.), *Inductive reasoning: Experimental, developmental, and computational approaches* (pp. 81-113). Cambridge, UK: Cambridge University Press.
- Spellman, B. A., & Holyoak, K. J. (1996). Pragmatics in analogical mapping. *Cognitive Psychology*, *31*, 307-346.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327-352.
- Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, *121*, 222-236.
- Winston, P. (1980). Learning and reasoning by analogy. *Communications of the ACM*, *23*, 689-703.