

Types of Cognitive Content and the Role of Relational Processing in the Illusion of Explanatory Depth

Graham Silk-Eglit (gsilke1@binghamton.edu)

Kenneth J. Kurtz (kkurtz@binghamton.edu)

Department of Psychology, Binghamton University,
Binghamton, New York 13902

Abstract

Rozenblit and Keil (2002) claim that people are subject to an illusion of explanatory depth (IOED) whereby they believe they understand the world in greater detail, coherence and depth than they actually do. In the present research, we questioned Rozenblit and Keil's conclusions in two ways. First, we tested whether people might overestimate their explanatory knowledge as a result of misconstruing how to initially rate their understanding of stimuli. We found that when directed to consider the physical-mechanical processes of stimuli instead of their functional affordances, participants did in fact offer more accurate estimates of understanding relative to their explanatory performance. Second, we tested whether the explanations participants proffer are misleadingly shallow. We predicted that by encouraging a more relational encoding of stimuli, participants would be able to produce better explanations. However, the results showed that participants' explanations remained shallow after relationally encoding stimuli.

Keywords: illusion of explanatory depth; analogical reasoning; modes of explanation; concepts and categories.

Introduction

Recent research has suggested that people often overestimate their ability to explain causally complex systems. In a recent study, Rozenblit and Keil (2002) asked participants to rate their understanding of devices both before and after proffering explanations of how those devices work. They found that people consistently overrated their understanding – displaying a significant drop between their pre-explanation and post-explanation ratings. Further analyses revealed that this overestimation was not a result of general overconfidence but was specific to causally complex systems; when tested with other, non-explanatory forms of knowledge—i.e., knowledge of facts, procedures, and narratives—this illusion of knowing was either absent (for procedures) or greatly diminished (for facts and narratives). Thus, Rozenblit and Keil argued that people are subject to an illusion of explanatory depth (IOED) wherein they have inflated metacognitive beliefs of their ability to explain causal phenomena yet possess only shallow understandings of the causal workings of those phenomena.

In the current study, we set out to question Rozenblit and Keil's (2002) argument in two ways. First, we intended to show that people were biased to overestimate their pre-explanation ratings of understanding. We believe this

occurred as a result of participants misconstruing the initial rating prompt as pertaining to a cognitively different set of content than was relevant to their proceeding explanatory task. Second, we sought to show that people had more explanatory knowledge than they demonstrated – causing them to underestimate their post-explanation knowledge. Specifically, we expected that participants would be able to articulate more causal knowledge if they were encouraged to draw comparisons between the stimuli and a relevantly similar second domain.

Types of cognitive content

One of the features that distinguishes knowledge of complex causal systems from other forms of knowledge is that complex causal systems are composed of distinct types of cognitive content. A device, for instance, can be described in terms of its functions, consequences and global appearance, or in terms of its components and the causal relations among them. Accordingly, while people may understand a given device's functional properties well, they may be less knowledgeable about its physical-mechanical processes. As such, judgments of understanding that confuse these two types of cognitive content are potentially susceptible to metacognitive miscalibrations. In this manner, an illusion of explanatory depth might arise from a discrepancy in the type of cognitive content that people use in order to make their judgments of understanding before and after proffering an explanation.

This, we claim, is precisely what occurs in Rozenblit and Keil's (2002) study (see also Alter, Oppenheimer & Zemla, 2010). In order to describe how a given object works, Rozenblit and Keil require participants to explain that object physical-mechanically in terms of its components and the causal relations among them. As such, post-explanation ratings, made in light of performance on this causal explanatory task, index participants' knowledge of physical-mechanical processes. However, Rozenblit and Keil fail to explicitly specify which sort of cognitive content their participants should estimate their initial, pre-explanation understandings of stimuli on – making their instructions potentially ambiguous. Rozenblit and Keil merely instruct participants “to rate on a 7-point scale how well you feel you understand each one” (p. 39). Consequently, the metacognitive miscalibration of the IOED may be an artifact of these ambiguous instructions. Specifically, Rozenblit and

Keil's instructions may mislead participants to construe the pre-explanation rating prompt as pertaining to the functional properties of a device when the relevant metacognitive assessment centers on that device's physical-mechanical processes.

Thus, in the current study, we asked one group of participants to make their pre-explanation ratings based on their general understanding of stimuli, following Rozenblit and Keil's (2002) procedure, and instructed another group to make their pre-explanation ratings based on their knowledge of the physical-mechanical workings of stimuli. We predicted that the latter group would exhibit better calibrated estimates of understanding as the content upon which they based their pre-explanation and post-explanation ratings was consistent.

Relational processing

Research on concepts and categorization has become increasingly interested in the role of explanatory knowledge in categorization. Previous research has documented the influence of explanatory knowledge on a number of categorization phenomena including: helping to define the features of new categories (Wisniewski and Medin, 1994), facilitating the acquisition of new categories (Pazzani, 1991), influencing categorization decisions (Lin & Murphy, 1997; Murphy & Medin, 1985), and guiding inductions generated about a given category (Ross & Murphy, 1999). These findings are linked to the *theory view* of conceptual representations which suggests that concepts are structured by intuitive theories that explain the appearance, function, and other aspects of exemplars. This approach stands in stark contrast to prototype and exemplar models of concepts that characterize conceptual representations as consisting of feature listings. Instead, the theory approach holds that intuitive theories relate (often causally) the features of concepts within structured representations. However, Keil (2005) has taken the illusion of explanatory depth as evidence that these intuitive theories are relatively sparse, which suggests that concepts may only be weakly structured.

Research on analogical reasoning has similarly placed great emphasis on the importance of relational structure in representations, but also emphasizes the role of structural alignment as a basis for knowledge change (Gentner, 1983). According to structure mapping theory, the theoretical inspiration for this manipulation, comparisons proceed by establishing structural alignments between two representations. This occurs through a mapping of the objects, attributes, and relations of the two representations in which correspondences between the two domains are established. People prefer structurally consistent mappings in which there is a *one-to-one correspondence* between the elements in the two representations and in which the arguments of corresponding predicates also correspond (*parallel connectivity*). Selection of possible interpretations of an analogy is further guided by the *systematicity*

principle, which states that a system of relations that is connected by higher-order constraining relations (such as causal relations) is preferred over an equal number of independent matches (Gentner & Colhoun, in press). Moreover, according to the *relational focus* assumption, relational matches are considered more important than object/attribute matches in the evaluation of analogical relatedness (Markman & Gentner, 2000; Gentner & Kurtz, 2006).

However, numerous studies have demonstrated that such deep systems of relations are seldom utilized by novices. Chi, Feltovich, & Glaser (1981), for instance, showed that novice representations tend to be weakly structured and primarily focused on surface features (i.e., the objects and attributes of representations); lacking in the structural relations and abstractions characteristic of the representations of experts. Moreover, Gentner, Ratterman, and Forbus (1993) demonstrated that people typically retrieve instances from memory on the basis of surface similarity, not the relational similarity required for deeper comprehension. Taken together, one might reasonably expect that the sort of information retrieved and represented by novices for the construction of explanations would similarly consist of mere surface features, lacking many of the deeper causal relations required for a sophisticated account.

In order to overcome these superficial retrieval and representation processes, researchers have utilized comparisons. Previous research has demonstrated that carrying out comparisons between two domains with differing surface features can promote *relational highlighting*, the bringing to light of common relational structure between two domains. Comparisons have been found to aid in overcoming difficulties in analogical transfer (Gick & Holyoak, 1983), and can lead to the development of deeper and more structural understandings of problems (Cummins, 1982) and causal systems (Kurtz, Miao, and Gentner, 2001).

In the current study, we sought to encourage the process of relational highlighting in order to increase the depth of participants' explanations. To do so, we proposed a novel means of inducing comparisons to motivate relational encodings. Specifically, we asked participants to generate a list of near category members, where a near category was defined as a category that shares a good deal in common with a specific object, but lacks some crucial feature(s). It was hoped that by carrying out this task, the core relational features of the stimulus would be highlighted. To help participants elaborate on this core content, they were also asked to place their near category members in a sentence frame that was intended to induce an alignment of both the similar and different relations between the near category and the stimulus. As we believe this sort of relational encoding facilitates access to and articulation of explanatory content, we predicted that carrying out this near category comparison

treatment prior to proffering an explanation would result in more structured and causally deeper explanations.

Outline of the present study

In the present study, we implemented a 2 X 2 factorial design with specificity of instructions and pre-explanation encoding prompt as between-subjects factors. Subjective judgments of understanding were recorded at various times throughout the experiment both before and after explanations were proffered. We predicted that pre-explanation ratings of understanding would be better calibrated when participants were instructed to base those ratings specifically on their knowledge of the physical-mechanical workings of stimuli rather than on their general knowledge of those stimuli. Additionally, we predicted that the post-explanation ratings would reflect greater depth of understanding for the relational encoding condition than for the control condition.

Method

Participants

Eighty-three undergraduate students from Binghamton University participated for course credit.

Materials and Design

Fifteen stimuli were used, all of which were names of object concepts. Thirteen stimuli served as distractors during the initial rating phase and two stimuli served as test stimuli to be explained by participants at a later phase in the experiment. Of the two test stimuli, one was selected from Rozenblit and Keil's (2002) set of test stimuli—i.e., a zipper—and the other was developed for the purposes of this study—i.e., a manual air pump.

All materials were presented to participants in packets. Participants received packets with different instructions depending on their condition. In all, there were four different conditions: Rozenblit and Keil's Original Instructions and the Control Encoding Task (OC), New Physical-Mechanical Instructions and the Control Encoding Task (PC), Original Instructions and the Relational Encoding Task (OR), and New Physical-Mechanical Instructions and the Relational Encoding Task (PR).

Procedure

Following Rozenblit and Keil's (2002) procedure, the full study consisted of seven phases. In phase 1, participants learned how to use a 7-point rating scale to indicate their understanding of stimuli by considering two training examples. The scale ranged from "1," which indicated a shallow understanding, to "7," which indicated a deep, expert understanding. Participants considered the same two training examples that were used in Rozenblit and Keil's study.

In phase 2 of the study, participants read a list of devices and rated their understanding of each one on the 7-point scale described during the training. Depending on their condition, some participants were told to rate how well they felt they generally understood each object in question as in Rozenblit and Keil's (2002) study (designated "O" in the condition acronyms). In contrast to Rozenblit and Keil's study, others were instructed to rate their understanding specifically on how deeply they understood the physical-mechanical workings of the devices (designated "P" in the condition acronyms).

In phase 3 of the study, a two stimulus subset was selected from the initial 15 stimuli and participants were asked to consider one of two questions – neither of which were included in Rozenblit and Keil's (2002) study. In the control encoding task condition, participants were asked to describe the appearance of a single instance of the stimulus category (designated "C" in the condition acronyms). In the relational encoding condition, participants were asked generate a list of near neighbor categories and then to place those near categories into a sentence frame that ensured explicit consideration of the relational similarities and differences between the two domains (designated "R" in the condition acronyms).

The rest of the procedure followed Rozenblit and Keil's (2002) study. In phase 4, participants were asked to write a detailed, step-by-step causal explanation of the two test stimuli. After participants provided an explanation, they were asked to re-rate their understanding of the two stimuli.

In phase 5 of the study, participants answered a diagnostic question for each stimulus that required critical knowledge about the mechanism. For example, participants were asked to explain how pulling the handle of a manual air pump caused air to enter the chamber. After participants answered the diagnostic question, they were again asked to rate how well they understood the phenomena.

In phase 6 of the study, participants read a brief expert explanation of the two stimuli and then re-rated their prior understanding of the stimuli in light of that description.

In phase 7 of the study, participants rated how well they understood the two stimuli after having read the expert explanation.

Results

Analyses focused on the difference in pre-explanation subjective ratings of understanding compared to post-explanation ratings. Based on Rozenblit and Keil's (2002) study, it was pre-decided that the rating following the diagnostic question—i.e., rating 3—most accurately reflected participants' post-explanation explanatory knowledge self-estimate. As such, a 2 (instructions: original vs. physical-mechanical) X 2 (encoding: control vs. relational) factorial ANOVA on the difference scores between rating 3 and rating 1 was implemented for the zipper, the manual air pump, and both stimuli combined.

When analyses collapsed across both stimuli, a main effect of instructions was found, $F(1,127) = 13.032, p = .000, d = .60$, suggesting that the physical-mechanical instructions significantly reduced the magnitude of the IOED. Follow-up simple effects showed that the physical-mechanical instructions significantly reduced the IOED between both the OC and PC conditions, $F(1,68) = 6.380, p = .014, d = -.61$, and the OR and PR conditions, $F(1,59) = 6.501, p = .013, d = -.65$, demonstrating the robustness of the physical-mechanical instructions manipulation. Additionally, a main effect of encoding was found, $F(1,127) = 8.178, p = .005, d = .45$, indicating that, contrary to our prediction, the relational encoding manipulation increased the magnitude of the IOED. Follow-up simple effects revealed that the relational encoding manipulation significantly increased the magnitude of the IOED between the OC and OR conditions, $F(1,56) = 4.511, p = .038, d = .55$. Additionally, the increase in the magnitude of the IOED between the PC and PR conditions approached significance, $F(1,71) = 3.465, p = .067, d = .43$.

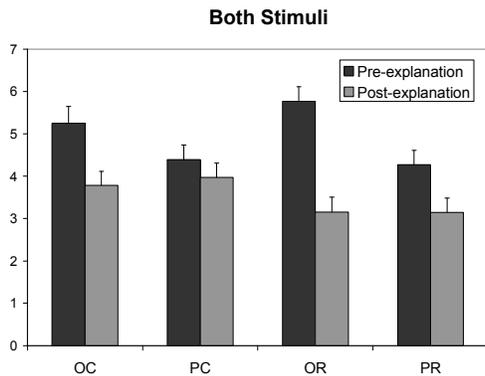


Figure 1: Mean pre- and post-explanation self-ratings of knowing in Rozenblit and Keil's Original Instructions and the Control Encoding Task (OC), New Physical-Mechanical Instructions and the Control Encoding Task (PC), Original Instructions and the Relational Encoding Task (OR), and New Physical-Mechanical Instructions and the Relational Encoding Task (PR) conditions for both stimuli.

For the zipper, analyses revealed a main effect of instructions, $F(1,61) = 5.728, p = .020, d = -.52$, suggesting that the physical-mechanical instructions significantly reduced the IOED. Follow-up simple effects demonstrated that although the difference scores for the OR condition compared to the difference scores for the PR condition approached significance, $t(1, 28) = -1/994, p = .056, d = -.75$, the difference scores for the OC compared to the PC condition did not, $t(1, 33) = -1.303, p = .202, d = -.44$. Additionally, a main effect of encoding was found, $F(1,61) = 7.627, p = .008, d = .61$. Again contrary to our prediction, this main effect suggested that the relational encoding manipulation increased the magnitude of the IOED. Follow-up simple effects showed that the participants in the OR condition exhibited a significantly larger IOED than those in

the OC condition, $t(1,26) = 2.369, p = .026, d = .90$; however, the difference scores of the PC condition relative to the PR condition were not significant, $t(1,35) = 1.523, p = .137, d = .50$.

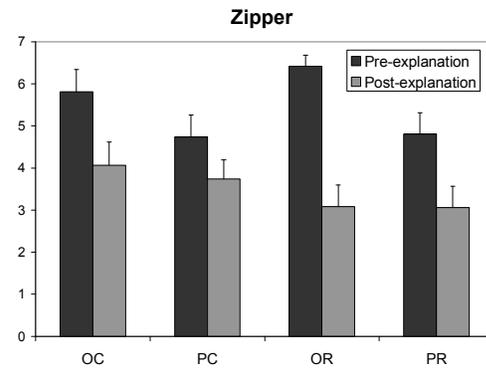


Figure 2: Mean pre- and post-explanation self-ratings of knowing in Rozenblit and Keil's Original Instructions and the Control Encoding Task (OC), New Physical-Mechanical Instructions and the Control Encoding Task (PC), Original Instructions and the Relational Encoding Task (OR), and New Physical-Mechanical Instructions and the Relational Encoding Task (PR) conditions for the zipper.

Analyses on the manual air pump again revealed a significant main effect of instructions, $F(1,62) = 9.191, p = .004, d = -.73$, suggesting that the physical-mechanical instructions significantly reduced the magnitude of the IOED. Follow-up simple effects demonstrated a significant reduction of the magnitude of IOEDs between the OC and PC conditions, $F(1,33) = 5.581, p = .024, d = .80$. Additionally, the reduction of the magnitude of IOEDs between the OR and PR conditions approached significance, $F(1,29) = 3.880, p = .058, d = .69$. However, a significant main effect of encoding was not found, $F(1,62) = 2.309, p = .134, d = .35$, indicating that the relational encoding task did not appear to increase the magnitude of the IOED for the manual air pump.

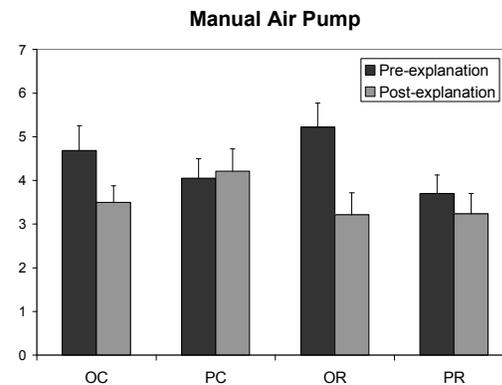


Figure 3: Mean pre- and post-explanation self-ratings of understanding in Rozenblit and Keil's Original Instructions and the Control Encoding Task (OC), New Physical-Mechanical Instructions and the Control Encoding Task

(PC), Original Instructions and the Relational Encoding Task (OR), and New Physical-Mechanical Instructions and the Relational Encoding Task (PR) conditions for the manual air pump.

Thus, our results confirmed our first prediction that specifying to participants that they base their pre-explanation ratings on their knowledge of the physical-mechanical workings of stimuli would reduce the magnitude of the IOED. Contrary to our second prediction, the relational encoding task appeared to increase the magnitude of the IOED especially for the zipper stimulus.

Discussion

The results of our experiment suggest that a confusion of relevant cognitive content on the initial, pre-explanation rating plays a significant role in the IOED. Specifically, our results show that when participants were explicitly encouraged to base their pre-explanation ratings of understanding on the physical-mechanical workings of stimulus devices, the IOED was eliminated for one stimuli (i.e., the manual air pump), significantly reduced for another (i.e., the zipper) and when both stimuli were combined the IOED was significantly reduced. In fact, the IOED was quite small when analyses collapsed across both stimuli. Consequently, we argue that Rozenblit and Keil's (2002) instructions permit participants to base their initial, pre-explanation ratings on a range of different types of cognitive content. Given the central role of functional features in artifact concepts (Bloom, 1996; Kelemen & Carey, 2007), it seems natural that participants would gauge their understanding of artifact concepts through a functional framework. The illusion of explanatory depth then appears in large part to index a discrepancy in the type of cognitive content used to determine one's depth of understanding, with pre-explanation ratings being based on knowledge of functional features and post-explanation ratings based on knowledge of physical-mechanical processes. As this study demonstrates, when this discrepancy in cognitive content is removed by specifying that pre-explanation ratings should be based exclusively on knowledge of physical-mechanical processes, the IOED is substantially reduced, if not eliminated.

Recent research on modes of explanation may offer some insight into this finding (Keil, 1994; Lombrozo, 2009). According to this research, people can adopt distinct modes of explanation when considering a given entity, including a physical-mechanical mode and a functional, or teleological, mode. While a physical-mechanical mode focuses on simple physical objects and their interactions, a functional mode centers on the goals or purposes of features and objects within an overall design. As such, modes of explanation posit different kinds of relations and properties as central and support different types of generalizations and predictions. For instance, to explain why a moving billiard ball stops on impact with a second billiard ball and to

predict what direction and at what speed the second billiard ball will travel, a physical-mechanical mode of explanation will do well. In contrast, to explain the rationale behind the configuration of keys on a computer's keyboard, a functional mode of explanation will do better. Importantly, Keil (1994), Atran (1995), and Kelemen (1999) have tied these modes of explanation to specific domains with functional modes of explanation applied to entities that exhibit actual or apparent design, especially artifacts. According to this account, an illusion of explanatory depth for physical-mechanical knowledge might arise from the assumption of a functional mode of explanation that misleads participants to base their level of understanding on certain relations and properties that are irrelevant to their self-evaluation of their physical-mechanical explanatory knowledge.

In contrast, the relational encoding manipulation failed to yield higher self-ratings of explanatory depth. This may have occurred for two reasons. First, the subjective basis of the post-explanation self-ratings may have been altered as a result of the relational encoding manipulation. This, in turn, could have happened for two reasons. On the one hand, by carrying out the relational encoding manipulation, participants may have realized the depth of the device being explained and may have judged their own explanation as more shallow as a result of that realization. In this sense, although the participants' explanations in the relational encoding conditions may have been deeper than those of their counterparts in the control encoding conditions, they may have rated them as more shallow as a result of having a greater understanding of the explanatory depth that they were not achieving. On the other hand, participants may have unwittingly made their post-explanation ratings based on the combination of the relational encoding/control encoding task and their actual explanation rather than simply on the basis of their explanatory performance alone. Given that the relational encoding task was intended to be more probing and therefore more difficult than the control encoding task, this may have led to lower self-ratings of understanding.

Secondly, the relational encoding task may have simply failed to elicit greater explanatory depth. This may have occurred because of participants' tendency to redundantly, and often time exclusively, offer lower-order relations in the relational encoding task. For example, for the zipper stimulus, participants frequently offered near neighbor categories that "hold things together," such as a button, Velcro, and a staple. However, very seldom did they expand on this lower-order relation to include the higher-order, causal relations that would explain how "holding things together" is achieved. It may be the case that these higher-order relations and the systematicity of the relations offered by participants are more important for eliciting greater causal explanatory knowledge than the amount of lower-order relations offered. A further study intended to elicit more higher-order, systematic relational encodings and to

incorporate objective, independent ratings of explanations is currently underway.

Limitations

One concern with the present study is that due to the small sample sizes in each condition, significant results were often found only when analyses collapsed across conditions or across stimuli. As such, many of our simple effect measures only approached significance. While the general pattern of findings was evident, our results may have been strengthened by a larger sample.

A second concern of the study, briefly addressed above, was the lack of inclusion of independent raters to objectively determine the depth of explanations. This leaves open the possibility that our relational encoding manipulation may have altered the subjective basis of participants' post-explanation self-ratings of understanding. Nonetheless, this possibility would not have affected our finding that specifying the type of content that participants should use to rate their pre-explanation understanding of stimuli effectively reduces, if not eliminates, the IOED.

Lastly, the findings of the present study—especially the reduction of the IOED resulting from the physical-mechanical pre-explanation instructions—should be replicated with additional stimuli. It would be best to replicate these findings not only with technological devices, but with other complex causal systems as well.

References

- Alter, A.L., Oppenheimer, D.M., & Zemla, J.C. (2010). Missing the trees for the forest: A construal level account of the illusion of explanatory depth. *Journal of Personality and Social Psychology*, 99(3), 436-451.
- Atran, S. (1995). Classifying nature across cultures. In D. Osherson and E. Smith (Eds.), *Invitation to Cognitive Science: Thinking* vol. 3, 2nd ed. MIT Press, Cambridge.
- Bloom, P. (1996). Intention, history, and artifact concepts. *Cognition*, 60, 1-29.
- Chi, M.T., Feltovich, P.J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5, 121-152.
- Cummins, D.D. (1992). Role of analogical reasoning in the induction of problem categories. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 18, 1103-1124.
- Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D., & Kurtz, K. (2006). Relations, objects, and the composition of analogies. *Cognitive Science*, 30, 609-642.
- Gentner, D., & Colhoun, J. (in press). Analogical processes in human thinking and learning. In A. von Müller & E. Pöppel Series Eds.) & B. Glatzeder, V. Goel, & A. von Müller (Vol. Eds.), *On Thinking: Vol. 2. Towards a Theory of Thinking*. Springer-Verlag Berlin Heidelberg.
- Gick, M.L., & Holyoak, K.J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 12, 306-355.
- Keil, F.C. (1994). The birth and nurturance of concepts by domains: The origins of concepts of living things. In L. Hirschfield and S. Gelman (Eds.), *Mapping the Mind: domain specificity in cognition and culture*. Cambridge University Press, New York.
- Keil, F.C. (2005). Knowledge, categorization, and the bliss of ignorance. In L. Gershkoff-Stowe, and D. Rakison (Eds.), *Building object categories in developmental time*. Hillsdale, NJ: Erlbaum.
- Kelemen, D. (1999). The scope of teleological thinking in preschool children. *Cognition*, 70, 241-272.
- Kelemen, D. & Carey, S. (2007). The essence of artifacts: Developing the design stance. In Laurence, S. & Margolis, E. (Eds.), *Creations of the mind: Artifacts and their representation*. Oxford University Press, Oxford.
- Kurtz, K.J., Miao, C.H., & Gentner, D. (2001). Learning by analogical bootstrapping. *Journal of Learning Sciences*, 10(4), 417-446.
- Lin, E.L., & Murphy, G.L. (1997). Effects of background knowledge on object categorization and part detection. *Journal of Experimental Psychology: Human Perception and Performance*, 23(4), 1153-1169.
- Lombrozo, T. (2009). Explanation and categorization: How “why?” informs “what?” *Cognition*, 110, 248-253.
- Markman, A.B., Gentner, D. (2000). Structure-mapping in the comparison process. *American Journal of Psychology*, 113(4), 501-538.
- Murphy, G.L., & Medin, D.L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.
- Pazzani, M.J. (1991). Influence of prior knowledge on concept acquisition: Experimental and computational results. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 416-432.
- Ross, B.H., & Murphy, G.L. (1999). Food for thought: Cross-classification and category organization in a complex real-world domain. *Cognitive Psychology*, 38, 495-553.
- Rozenblit and Keil (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, 92, 1-42.
- Wisniewski, E.J., & Medin, D.L. (1994). On the interaction of theory and data in concept learning. *Cognitive Science*, 18, 221-281.