



Cognitive Science 36 (2012) 919–932

Copyright © 2012 Cognitive Science Society, Inc. All rights reserved.

ISSN: 0364-0213 print / 1551-6709 online

DOI: 10.1111/j.1551-6709.2012.01253.x

Causal Systems Categories: Differences in Novice and Expert Categorization of Causal Phenomena

Benjamin M. Rottman,^a Dedre Gentner,^b Micah B. Goldwater^b

^a*Department of Hospital Medicine, University of Chicago*

^b*Department of Psychology, Northwestern University*

Received 24 May 2011; received in revised form 7 October 2011; accepted 23 October 2011

Abstract

We investigated the understanding of causal systems categories—categories defined by common causal structure rather than by common domain content—among college students. We asked students who were either novices or experts in the physical sciences to sort descriptions of real-world phenomena that varied in their causal structure (e.g., negative feedback vs. causal chain) and in their content domain (e.g., economics vs. biology). Our hypothesis was that there would be a shift from domain-based sorting to causal sorting with increasing expertise in the relevant domains. This prediction was borne out: The novice groups sorted primarily by domain and the expert group sorted by causal category. These results suggest that science training facilitates insight about causal structures.

Keywords: Causal reasoning; Categorization; Relational categories; Expertise

Causality is of central importance in human cognition. We employ causal reasoning in explanation and prediction, in category organization, and in goal-based planning. For this reason, causal knowledge and reasoning has been a focus of cognitive science from the outset (Forbus, 1984; Hayes, 1979; de Kleer & Brown, 1981). Recent research on causal reasoning has sought to capture the details of how people think about causality, with a variety of methodological approaches (Sloman, 2005). Some work focuses on capturing relations between variables in graphical representations such as qualitative process models (Forbus, 1985; Forbus, Nielsen, & Faltings, 1991) or causal Bayesian networks (e.g., Gopnik et al., 2004; Waldmann, Hagmayer, & Blaisdell, 2006), while other work focuses on characterizing the kinds of experiences that lead people to infer a causal relationship, including particular kinds of statistical relations among variables (Cheng, 2000) and evidence of causal mechanisms (Ahn & Kalish, 2000; Ahn, Kalish, Medin, & Gelman 1995; Friedman & Forbus, 2008). Researchers have also investigated the relation between causal

Correspondence should be sent to Dedre Gentner, Psychology Department, Northwestern University, 2029 Sheridan Road, Evanston, IL 60208. E-mail: gentner@northwestern.edu

beliefs and explanation (Lombrozo & Carey, 2006), between causation and category structure (Ahn, Kim, Lassaline, & Dennis, 2000; Rehder & Burnett, 2005), and between causal models and language (Kuehne & Forbus, 2002; Wolff, 2003; Wolff & Song, 2003). In general, this research has focused on how people learn and reason about a particular causal structure, such as learning that asbestos leads to DNA mutation, which leads to cancer (e.g., Ahn et al., 2000; Lagnado, Waldmann, Hagmayer, & Sloman, 2007; Rehder & Hastie, 2001; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003).

Our focus here is different. We ask whether and to what extent people possess an abstract understanding of causal structure that allows recognition of common causal structures across disparate domains. For example, does understanding one positive feedback system—say audio feedback in an acoustic system—allow recognition of positive feedback in a different domain, such as a pricing bubble in economics? Or to put it another way, what sort of experience is necessary in order for someone to see these two phenomena as sharing important causal structure?

In this paper, we examine the representation and use of *causal system categories*, of which *positive feedback systems* is an example. Causal systems categories are defined by possessing common causal structure, irrespective of the particular domain. For example, the phenomena of population growth, economic bubbles, electronic audio feedback, and melting polar ice caps are all governed by *positive feedback*: a causal structure in which (for example) an event X positively influences another event Y, which in turn positively influences X, producing a cycle of increasing magnitude of effect in which the output becomes more and more extreme. There may be many more events than just two, but the principle is the same. Another causal category is a *causal chain*, in which an event X influences Y, which in turn influences Z.

Although prior research has investigated how people learn and reason about particular causal structures such as causal chains or common cause structures (e.g., Fernbach & Sloman, 2009; Kim, Luhmann, Pierce, & Ryan, 2009; Rehder, 2003a,b; Rehder & Burnett, 2005; Waldmann, 2000), it has not addressed whether people mainly treat each phenomenon as an individual, or whether they recognize that the phenomena form a class based on common causal structure. We consider three possibilities.

First, given the centrality of causal relations in human cognition, it may be that adults in general have an abstract understanding of causal system commonalities, and that they view multiple phenomena with the same structure as comprising a class. However, a second, equally plausible alternative is that people focus on causal patterns within a given domain and fail to perceive general patterns of causal structure across domains. After all, the causal learning literature shows that even seemingly simple causal structures like causal chains can be quite difficult to learn (Steyvers et al., 2003). Given the complexity of many causal phenomena—such as the feedback structures involved in global warming or diabetes—it could well be that people's causal reasoning remains chiefly at a domain-specific level.

The third possibility is that people begin with domain-specific causal models, but that with increasing education and experience, people can come to see abstract domain-general causal patterns. This third possibility gains credence from research on conceptual development. There is considerable evidence that relational information is acquired later than

concrete featural information. For example, when asked to interpret the comparison “A cloud is like a sponge,” 4–5-year-old children focus on common concrete properties, such as “both are round and fluffy” while adults and older children focus on relational commonalities, such as “both store water and then give it back to you” (Gentner, 1988). In the same vein, when asked to reenact a story with new characters, 6-year-olds rely heavily on perceptual similarities among the characters, whereas 9-year-olds can map the plot structure even to very different characters (Gentner & Toupin, 1986).

In adults, a similar relational shift pattern also occurs with increases in domain knowledge. For example, Shafto and Coley (2003) found that while novices sorted fish by similarity of appearance, commercial fisherman grouped them according to behavioral and causal/ecological relations (see also, Medin et al., 2006; Proffitt, Coley, & Medin, 2000). Chi, Feltovich, and Glaser (1981) found that when expert and novice physics students were asked to sort physics problems, experts sorted the problems according to common principles (e.g., conservation of energy), which define the abstract structure of the problem, whereas novices sorted based on concrete similarities (e.g., problems containing ramps). Likewise, novice chemistry students classify chemical reactions in terms of concrete features of chemical reactions, such as whether a reaction produces water, whereas experts classify reactions by common chemical mechanisms, such as acid–base reactions (Stains & Talanquer, 2008). Analogously, we predicted that a shift from focusing on concrete domain properties to the more abstract causal relational structures would occur in the way people organize causal phenomena.

In the current experiment, we examined five causal system categories (see Fig. 1). The first is a *common effect* structure, a causal structure in which many factors influence one effect. The level of CO₂ in the atmosphere is a product of fossil fuel burning, plant photosynthesis, and CO₂ absorption in oceans. Heart disease is caused by genetics, smoking, high blood cholesterol, high blood pressure, and obesity. A *common cause* structure is one in which a single cause has multiple effects. For example, an allergic reaction can cause multiple effects, including inflammation, rash, asthmatic reactions, and sneezing. Another






	A: Common Effect	B: Common Cause	C: Positive Feedback	D: Negative Feedback	E: Chain
					
1: Environmental Science					
2: Economics					
3: Electrical Engineering					
4: Mechanical Engineering					
5: Biology					

Fig. 1. Matrix of materials used in the study (5 causal categories × 5 domains). Cells with gray backgrounds designate the phenomena that served as seed cards.

example, from economics, is that the unemployment rate has causal effects on a number of phenomena, including the crime rate, suicide rate, general health conditions, and GDP.

The third causal structure, defined above, is a *positive feedback* structure—often described as a “vicious cycle.” In a positive feedback structure, the output is “fed back” in such a way as to magnify the input. This in turn produces a greater output, resulting in a cycle of increasing magnitude of effect. For example, increased global temperatures cause polar ice to melt. Whereas ice reflects sunlight well, water absorbs more sunlight, which means that the earth warms faster. This in turn causes more polar ice to melt. Economic bubbles work in a similar way. People buy stocks or property because they assume the prices will increase. The buying itself causes the prices to increase, which causes yet more buying, and so on. *Negative feedback* is a causal structure in which the output of the structure is fed back in such a way as to reduce the input. If the input is increased, the resulting output will also increase, and this will act to reduce the subsequent input. This results in a cycle that stabilizes the system. For example, in the domain of biology, humans regulate temperature by perspiring when they are too warm. When the body has cooled enough, it stops perspiring. Similarly, in the domain of economics, the Federal Reserve regulates the economy by raising, lowering, or keeping constant the interest rate. If the economy is slow, the FED will lower interest rates to stimulate borrowing and economic activity. If the economy is growing too fast, the FED will raise interest rates to encourage saving.

The last structure is a *causal chain*—a structure wherein an event A influences an event B, which in turn influences C, etc. The chain can be of any length; the sequential nature is the chain’s defining characteristic. For example, we can understand the way petroleum prices affect the price of consumer goods as a causal chain: the price of gas affects the price of transportation of goods, which affects the price of goods sold in stores. Another example of a causal chain, from neurobiology, is that synaptic transmission requires a series of steps, including an action potential, release of neurotransmitter, and binding on the postsynaptic neuron.

Our interest here is in whether people organize their knowledge of causal phenomena according to these kinds of causal categories. To investigate this, we set up a design in which people could sort phenomena such as those described above either by domain content or by causal category. We created 25 descriptions of phenomena, organized into a 5×5 matrix defined by crossing the five causal categories just described with five content domains: biology, environmental science, economics, electrical engineering, and mechanical engineering (Fig. 1; see online supplement for full stimuli). Then we asked college students to sort the phenomena into categories, as described below. Our questions are (1) whether people will chose to sort by domain or by causal structure; and (2) whether their preferred organization will be influenced by expertise.

The participants were drawn from four populations of Northwestern University students differing in their expertise in *physical science*. Three groups were students majoring in different social sciences: economics, sociology, and psychology. The fourth group was made up of students majoring in the physical sciences, primarily physics majors and students in the Integrated Science Program. Many of the physical science students, particularly those in the Integrated Science Program, had been trained in multiple sciences. Thus, in

addition to broader knowledge of the physical sciences, they had another possible advantage: Their course work in different physical sciences could have encouraged them to compare and abstract across domains.

For our four populations, we have two predictions and an open question. The psychology and sociology students were considered to be novices for purposes of the study, as their expertise does not match any of the five domains of phenomena used in the stimuli. We predicted these students would sort chiefly by content domain. At the other extreme, the physical science students served as the expert group. Because of their training in multiple relevant content domains, we expected that they would sort largely by causal category. The economics majors were chosen because they have experience in one of the content domains used in the study. If expertise in one domain is sufficient to recognize and use cross-domain commonalities in organizing causal phenomena, then the economics majors should show more causal-category sorting than the psychology and sociology majors.

1. Methods

1.1. Participants

Forty-four undergraduate students from Northwestern University were recruited from four populations: introductory psychology students (12), sociology majors (9), economics majors (11), and physical science students (primarily physics and Integrated Science Program majors) (12). The psychology students participated in exchange for course credit. The other three populations were compensated for their participation.¹

1.2. Materials calibration

The 25 phenomena descriptions varied orthogonally on two dimensions, causal system and content domain (see Fig. 1 and Table 1). By the nature of the design, passages coming from the same domain should share more semantic similarities than did those from the same causal system. To confirm this, we use Latent Semantic Analysis (LSA)—a mathematical method for inducing the degree of relatedness between words in a large body of text on the basis of their contextual co-occurrence patterns—to generate pair-wise relatedness scores for the descriptions (Landauer & Kintsch, 2006).

For each of the 25 phenomena, we calculated the mean LSA relatedness ratings to other cards within the same domain ($M = 0.21$, $SD = 0.08$), within the same causal system category ($M = 0.10$, $SD = 0.04$), and with neither of those relations ($M = 0.08$, $SD = 0.02$). Pairwise t -tests revealed that pairs of descriptions from the same domain had significantly higher relatedness ratings than pairs that shared the same causal system, $t(24) = 5.59$, $p < .01$, and pairs that shared neither $t(24) = 7.22$, $p < .01$. Additionally, the within causal system pairs had marginally higher relatedness ratings than the no relation pairs, $t(24) = 2.10$, $p = .05$. Thus, as intended, the overall semantic relatedness is high for items from the same domain, but not for items that share causal structure.

Table 1

Examples of two seed cards, each showing one domain match and one causal systems match

Seed Card: (Environmental Science, Common Effect)	
Many processes are responsible for the level of CO ₂ in the atmosphere such as fossil fuel burning, plant photo synthesis, and CO ₂ absorption in oceans. When making public policy decisions about environmental issues, we must keep in mind the many complex factors that influence the level of CO ₂ in the atmosphere.	
Same-Domain Match: (Environmental Sci., Positive Feedback)	Same Causal System Match: (Biology, Common Effect)
In the process of global warming, as the temperature of the earth rises, polar ice begins to melt. Water absorbs more heat from sunlight than ice. Consequently, as ice is turned into water, the temperature of the earth begins to rise even faster, which in turn leads to increased ice-melt.	Aside from genetics, the four main factors which increase the risk of heart disease are smoking, high blood cholesterol, high blood pressure, and obesity. Doctors advise people with any one, or a combination of these risk factors to improve their health to mitigate the risk of heart attack.
Seed Card: (Electrical Engineering, Negative Feedback)	
A thermostat works by measuring temperature and turning on or off a furnace or air conditioner to reach a desired temperature. If the temperature is too cold, the thermostat will turn on the furnace until it becomes warm enough. Likewise, the thermostat on an air conditioner turns on when the house is too warm.	
Same-Domain Match: (Elect. Engineering, Common Cause)	Same Causal System Match: (Economics, Negative Feedback)
Internet routers work by distributing a data signal to multiple devices. If the router is turned off, then all the computers lose signal. However, the functioning of one individual computer does not affect the functioning of other computers. Thus, if one computer is turned off, all the others still get the data signal from the router.	The Federal Reserve has the ability to raise or lower the interest rates depending on the current state of the economy. If the economy is slow, the FED will lower interest rates to stimulate borrowing and economic activity. Raising interest rates will slow the economy by increasing the cost of borrowing.

1.3. Procedure

Each of the 25 phenomena was printed on a file card, and five cards (the grayed cells in Fig. 1) were designated as “seed cards.” Participants were told that they were to sort descriptions of real-world phenomena into categories that “go together.” On the table before them there were five columns headed by the seed cards, plus one column labeled “Other.” Participants were asked first to read the seed cards, and then to sort the remaining 20 cards into the columns “based on how well the description on the card goes with the initial card” already in the column. They could also place a card into the “Other” column if they did not think it fitted with any of the seed cards.

Because our goal was to pit the two sorting strategies (content domain and causal system) against one another, the seed cards were chosen to be equally applicable to either strategy. Specifically, the five seed cards each differed from one another both in content domain and in causal category. Thus, the columns could be taken to represent five different domains or five different causal systems, or some mixture of the two.

The 20 cards to be sorted were arranged in a semi-random order such that two cards from the same domain were never in immediate sequence, with half the participants receiving the

reverse order. There was no time limit; participants were told that they could work with the cards in any order and were allowed to rearrange previously sorted cards.

2. Results

Each of the 20 cards (one “sorting”) was coded according to whether it matched the seed card in its column by content domain or by causal system. On average, across all groups, 5.75 ($SD = 2.61$) of the 20 cards either were placed in the “other” column or failed to match their seed card by either domain or causal system. This average was consistent for all four groups and there was no main effect of group in a one-way ANOVA, $F(3,44) < 1$. Since we were primarily interested in the cards sorted by domain or causal system, these cards were ignored in the following analyses.

Out of the cards that a given participant sorted either by domain or by causal structure, we looked at the percentage that were sorted by causal structure. On average, the physical science students sorted 61% of these cards ($SD = 39\%$) by causal structure. However, for the other two groups, the percentage of cards sorted by causal structure was smaller: $M = 34\%$, $SD = 33\%$ for the psychology and sociology students, $M = 36\%$, $SD = 33\%$ for the economics students. A t -test did not find a difference between the group of psychology and sociology students versus the economics students, $t(30) < 1$. Therefore, we considered psychology, sociology, and economics students as a group (*social science* students) for further analyses. Importantly the physical science students sorted significantly more cards by causal structure than the combined social science group, $M = 35\%$, $SD = 32\%$, $t(42) = 2.25$, $p = .03$

To better understand the dominant sorting strategy used by participants of different groups, hierarchical cluster analysis (HCA) was performed based on a pairwise co-occurrence matrix per group (how many times participants sorted each pair of cards together). The HCA used the Euclidean squared distance metric and the between groups linkage agglomeration method. One benefit of the HCA analysis compared to the previous analysis was that the previous analysis only identified whether a card was sorted with a same-domain seed or with a same-causal-system seed. The HCA analysis looks at all of the relationships between the 20 sorted cards, ignoring the seed cards.

For the physical science students, HCA revealed five identifiable clusters of phenomena, closely matching the five causal system categories (Fig. 2; cards with the same letter of the alphabet share the same causal structure category). Three cards were not generally sorted by causal structure. The card E1 (a chain-environmental science card) was generally sorted in the positive feedback group (C), particularly with the card C1 (a positive feedback-environmental science card). Both of these cards were about global warming. Likewise, A4 was sorted with E4, and E2 with A2. Thus, it appears that high domain relatedness may have overridden the dominant strategy of sorting by causal systems. In contrast, the HCA dendrograms for the social science students (Fig. 3; cards with the same number share the same domain), revealed that the social science students sorted primarily by common domain, with the domains of mechanical and electrical engineering somewhat conflated.

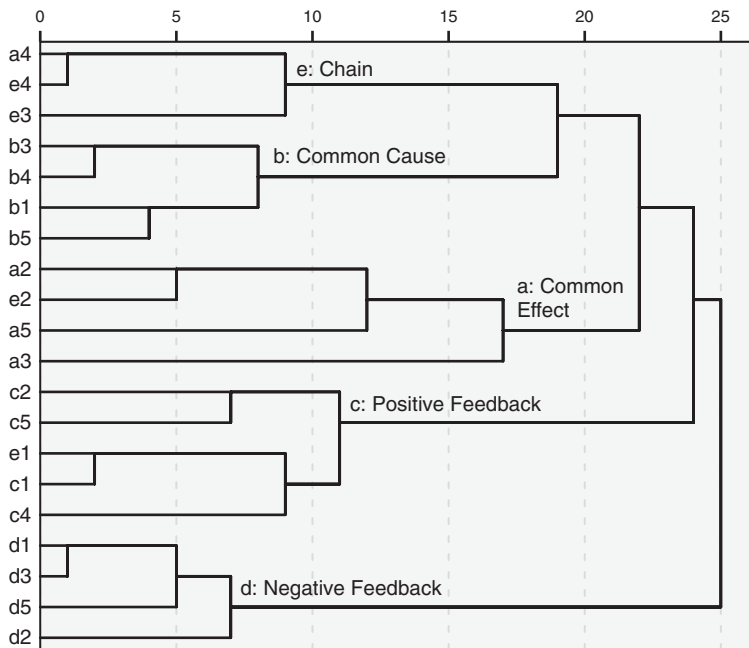


Fig. 2. Hierarchical cluster analysis for the physical science group. *Note:* Letters refer to the causal structures groups and numbers refer to the content domains (see Fig. 1).

To further probe into the basis for students' sorting decisions, we computed the LSA relatedness scores for the co-sorted cards. The mean LSA score for the co-sorted clusters for the combined social science group ($M = .17$, $SD = .04$) was nonsignificantly higher than that of the physical science group ($M = .14$, $SD = .05$), $t(42) = 1.93$, $p = .06$. It is possible that the social science students relied in part on local semantic connections between terms from the same content domain.

3. Discussion

The same patterns of causation occur across different domains. Our question is whether and to what degree people perceive these abstract causal structures. Specifically, we asked whether there are novice-expert differences in the degree to which students perceive the common causal structure underlying different phenomena, as assessed by their sorting behavior in a task that allows sorting either by causal system or by content domain. We predicted that experts (physical science students) would show more causal sorting than novices (psychology and sociology students). This prediction was confirmed: Psychology and sociology students sorted primarily by content domain, and physical science students sorted primarily by causal system. Additionally, we asked whether economics students would pattern with the other social science students or with the physical science students. They

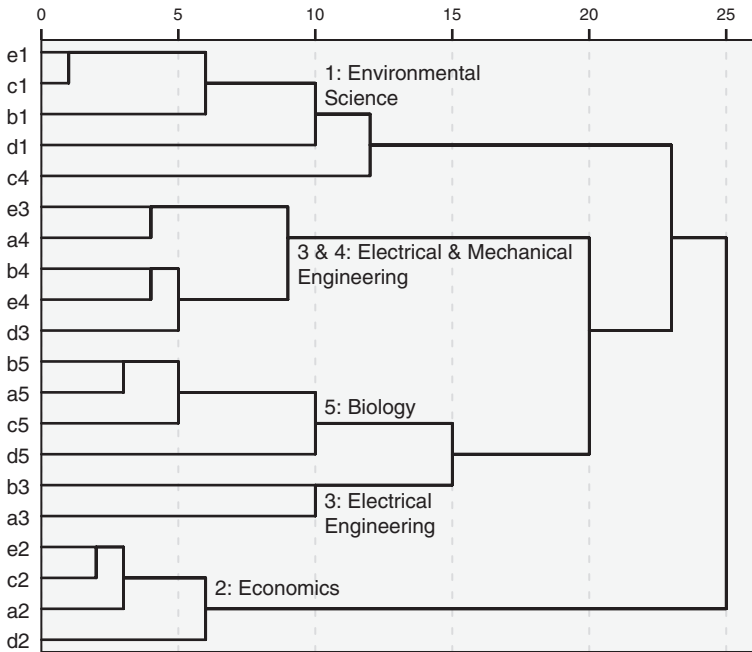


Fig. 3. Hierarchical cluster analysis for the social science group. *Note:* Letters refer to the causal structures groups and numbers refer to the content domains (see Fig. 1).

behaved as the other social science students did, sorting primarily by content domain. Hierarchical Cluster Analysis provided additional support for the hypothesis: HCA revealed clusters based on the domains for the combined social science group (including economics), and clusters based on causal systems for the physical science students.

Having an abstract understanding of causality, as described here, might confer a number of advantages. Previous research (Steyvers, Tenenbaum, Wagenmakers, & Blum 2003) has found significant individual differences in peoples’ ability to learn causal structures in a bottom-up statistical fashion. It is possible that some of these differences stem from differences in people’s abstract understanding of how different classes of causal structures function (see also Fernbach & Sloman, 2009). Another advantage of having an abstract understanding of causal structures is that it might help people interpret scientific evidence. For example, after learning that two variables are correlated, people with an abstract understanding of how different types of causal structures can produce a correlation between two variables might be more likely to search for a common cause instead of simply concluding that there must be a direct causal relation between the two variables.

Why did the physical science students, but not the social science students, sort causally? The most obvious explanation, as noted in the introduction, rests on differences in domain knowledge—specifically, in differences in the number of domains with which students were familiar. The physical science students typically had training in four of the five relevant domains—mechanical engineering, biology, electrical engineering, and environmental

science. In contrast, economics students had training in only one of the domains (economics), and psychology/sociology students in none of them. Domain experience might plausibly lead to a better articulation of general causal patterns within the domain. So perhaps the difference results simply from knowledge of more individual domains.

Another possibility is that social science students could have perceived the causal commonalities, but simply considered the domain-level commonalities more important. After all, reasoning well within a domain is important, and nothing about the task communicated that domain sorting was “wrong.” However, a follow-up study (Goldwater & Gentner, unpublished data) suggests that there were real differences in the perception of causal systems. In this study, social science students were given two opportunities to sort and were asked to use a different strategy for the second. Consistent with the current findings, the majority of the students were domain sorters, but some sorted causally to a fair degree. These better causal sorters went on to sort quite well by domain for their second sort. In contrast, the domain sorters did not adopt a causal strategy for their second sort. This suggests that the ability to perceive domain-level commonalities is widely available, but that perceiving common causal systems is not. While this does not mean that causal sorting is necessarily the “right” answer, it does suggest it is a skill that does not come for free.

What led the physical science students to show greater awareness of causal systems? We speculate that this difference was not simply a matter of greater domain expertise in multiple disciplines. Rather, it seems likely that the physical science students (particularly the Integrated Science students) had also acquired a stock of cross-domain abstractions. That is, their multidisciplinary experience had provided them with opportunities to compare across different domains and to extract general patterns such as positive feedback and common-cause. These abstractions might then leap readily to mind as a natural way to organize phenomena in a sorting task. One indication that this might be the case is that the physical science students were more likely to produce causal sorts of economics phenomena than were the other groups, including the economists. Out of 4 economics cards to sort, the physical science students sorted on average 1.37 of them causally, in contrast to only 0.73 for the economics students (and 0.78 for the combined social science group). This suggests that the physical science students had formed abstract causal categories and were able to classify new phenomena from economics into those categories.

This speculation that cross-domain comparison experience might support the formation of causal systems categories gains credence from prior work on the development and learning of relational categories (Gentner, Anggoro & Klibanoff, 2011; Goldwater & Markman, 2011; Kotovsky & Gentner, 1996). *Relational categories* are categories such as *catalyst* and *solvent*, whose membership is determined by common relational structure; they contrast with *entity categories*, which can be defined by common intrinsic properties, such as *beaker* and *pipette* (Asmuth & Gentner, 2005, unpublished data; Barr & Caplan, 1987; Gentner & Asmuth, 2008; Gentner & Kurtz, 2005; Goldwater, Markman, & Stilwell, 2011; Goldwater & Markman, 2011; Markman & Stilwell, 2001; Rein, Goldwater, & Markman, 2010; Ross & Murphy, 1999). The causal systems categories considered here clearly qualify as relational categories.

Developmental research shows that relational categories are slow to be acquired, relative to entity categories (Gentner, 2005). Further, earlier stages of learning are characterized by a focus on within-domain surface similarities; the ability to recognize relational structures across different situations emerges later (see Gentner & Rattermann, 1991; Doumas, Hummel, & Sandhofer, 2008 for reviews). There is also considerable evidence that analogical comparison promotes learning relational categories (Christie & Gentner, 2010; Gentner, Anggoro, & Klibanoff, 2011; Goldwater & Markman, 2011). Thus, one explanation for the difference between physical science students and the other groups is that the physical science students had many opportunities to compare across domains and had thereby abstracted a rich stock of causal abstractions.

An understanding of causal systems is crucial to explaining and predicting complex phenomena, both in the natural world and in social and economic spheres. The present work shows large differences in the ability to perceive common causal patterns across domains. Future work should reveal how people come to learn these causal patterns and to perceive abstract causal systems categories.

Note

1. No demographics were retained for the introductory psychology students. Out of 9 sociology majors, 6 were double-majors in other fields, including political science (2), communication studies, legal studies, psychology, and a major called Mathematical Methods in the Social Sciences. Out of the 11 economics majors, 7 were double majors in other fields, including math, music, international studies, electrical, biomedical, and industrial engineering, and Mathematical Methods in the Social Sciences. Out of the 12 physical science students, 4 were obtaining triple majors, 5 were double majors, and 3 were obtaining one major. Their majors were in the following fields: inter-disciplinary Science Program (8), physics (6), math (5), chemistry (3), biology (2), and classics (1).

Acknowledgments

This research was supported by Office of Naval Research grant N00014-92-J-1098 awarded to the second author. We are grateful to the Similarity and Analogy group at Northwestern University for many discussions of these issues, and to Kathleen Braun for help with the research and analyses.

References

- Ahn, W. K., & Kalish, C. W. (2000). The role of mechanism beliefs in causal reasoning. In F. C. Keil & R. A. Wilson (Eds.) *Explanation and cognition* (pp. 199–225). Cambridge, MA: MIT Press.

- Ahn, W. K., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation vs. mechanism information in causal attribution. *Cognition*, *54*, 299–352.
- Ahn, W., Kim, N. S., Lassaline, M. E., & Dennis, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology*, *41*, 361–416.
- Asmuth, J., & Gentner, D. (2005). Context sensitivity of relational nouns. In B. G. Bara, L. Barsalou, & M. Bucciarelli (Eds.), *Proceedings of the twenty-seventh annual meeting of the Cognitive Science Society* (pp. 163–168). Mahwah, NJ: Lawrence Erlbaum Associates.
- Barr, R. A., & Caplan, L. J. (1987). Category representations and their implications for category structure. *Memory and Cognition*, *15*, 397–418.
- Cheng, P. W. (2000). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.
- Chi, M., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, *5*, 121–152.
- Christie, S., & Gentner, D. (2010). Where hypotheses come from: Learning new relations by structural alignment. *Journal of Cognition and Development*, *11* (3), 356–373.
- Doumas, L. A. A., Hummel, J. E., & Sandhofer, C. M. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review*, *115*, 1–43.
- Fernbach, P. M., & Sloman, S. A. (2009). Causal learning with local computations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 678–693.
- Forbus, K. D. (1984). Qualitative process theory. *Journal of Artificial Intelligence*, *24*, 85–168.
- Forbus, K. D. (1985). The role of qualitative dynamics in naive physics. In J. Hobbs & R. Moore (Eds.), *Formal theories of the commonsense world* (pp. 185–226). Palo Alto, CA: Tioga Press.
- Forbus, K., Nielsen, P., & Faltings, B. (1991). Qualitative spatial reasoning: The CLOCK project. *Artificial Intelligence*, *51* (1–3), 417–471.
- Friedman, S., & Forbus, K. (2008). Learning causal models via progressive alignment and qualitative modeling: A simulation. In B. C. Love, K. McRae & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 1123–1128). Washington, DC: Cognitive Science Society.
- Gentner, D. (1988). Metaphor as structure mapping: The relational shift. *Child Development*, *59*, 47–59.
- Gentner, D. (2005). The development of relational category knowledge. In L. Gershkoff-Stowe & D. H. Rakison, (Eds.), *Building object categories in developmental time* (pp. 245–275). Hillsdale, NJ: Erlbaum.
- Gentner, D., Anggoro, F. K., & Klibanoff, R. S. (2011). Structure-mapping and relational language support children's learning of relational categories. *Child Development*, *82*, 1173–1188.
- Gentner, D., & Asmuth, J. (2008). Can relationality be distinguished from abstractness in noun mutability? In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th annual conference of the Cognitive Science Society* (pp. 863–868). Austin, TX: Cognitive Science Society.
- Gentner, D., & Kurtz, K. (2005). 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* (pp. 151–175). Washington, DC: APA.
- Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes (Eds.), *Perspectives on thought and language: Interrelations in development* (pp. 225–27). London: Cambridge University Press.
- Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, *10*, 277–300.
- Goldwater, M. B., & Markman, A. B. (2011). Categorizing entities by common role. *Psychonomic Bulletin and Review*, *18* (2), 406–413.
- Goldwater, M. B., Markman, A. B., & Stilwell, C. H. (2011). The empirical case for role-governed categories. *Cognition*, *118*, 359–376.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 3–32.
- Hayes, P. J. (1979). The naive physics manifesto. In D. Michie (Ed.), *Expert systems in the micro-electronic age* (pp. 242–270). Edinburgh, Scotland: Edinburgh University Press.

- Kim, N. S., Luhmann, C. C., Pierce, M. L., & Ryan, M. M. (2009). The conceptual centrality of causal cycles. *Memory & Cognition*, 37, 744–758.
- de Kleer, J., & Brown, J. S. (1981). Mental models of physical mechanisms and their acquisition. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 285–309). Hillsdale, NJ: Erlbaum.
- Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, 67, 2797–2822.
- Kuehne, S. E., & Forbus, K. D. (2002). Qualitative physics as a component in natural language semantics: A progress report. In W. D. Gray & C. D. Schunn (Eds.), *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society* (pp. 554–560). Fairfax, VA: George Mason University.
- Lagnado, D. A., Waldmann, M. R., Hagmayer, Y., & Sloman, S. A. (2007). Beyond covariation: Cues to causal structure. In A. Gopnik & L. Schultz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 154–172). Oxford, England: Oxford University Press.
- Landauer, T., & Kintsch, W. (2006). Latent Semantic Analysis. [Online Computer Software]. Accessed January 27, 2006, from <http://lsa.colorado.edu/>
- Lombrozo, T., & Carey, S. (2006). Functional explanation and the function of explanation. *Cognition*, 99 (2), 167–204.
- Markman, A. B., & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental and Theoretical Artificial Intelligence*, 13 (4), 329–358.
- Medin, D. L., Ross, N., Atran, S., Cox, D., Coley, J., Proffitt, J., & Blok, S. (2006). Folkbiology of freshwater fish. *Cognition*, 99 (3), 237–273.
- Proffitt, J. B., Coley, J. D., & Medin, D. L. (2000). Expertise and category-based induction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26 (4), 811–828.
- Rehder, B. (2003a). A causal-model theory of conceptual representation and categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 1141–59.
- Rehder, B. (2003b). Categorization as causal reasoning. *Cognitive Science*, 27, 709–748.
- Rehder, B., & Burnett, R. C. (2005). Feature inference and the causal structure of categories. *Cognitive Psychology*, 50, 264–314.
- Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: The effects of causal beliefs on categorization, induction, and similarity. *Journal of Experimental Psychology: General*, 130, 323–360.
- Rein, J. R., Goldwater, M. B., & Markman, A. B. (2010). What is typical about the typicality effect in category-based induction? *Memory and Cognition*, 38 (3), 377–38.
- 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.
- Shafto, P., & Coley, J. D. (2003). Development of categorization and reasoning in the natural world: Novices to experts, naïve similarity to ecological knowledge. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 29, 641–649.
- Sloman, S. A. (2005). *Causal models: How people think about the world and its alternatives*. Oxford, England: Oxford University Press.
- Stains, M., & Talanquer, V. (2008). Classification of chemical reactions: stages of expertise. *Journal of Research in Science Training*, 45, 771–793.
- Steyvers, M., Tenenbaum, J., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27, 453–489.
- Waldmann, M. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 53–76.
- Waldmann, M. R., Hagmayer, Y., & Blaisdell, A. P. (2006). Beyond the information given: Causal models in learning and reasoning. *Current Directions in Psychological Science*, 15 (6), 307–311.
- Wolff, P. (2003). Direct causation in the linguistic coding and individuation of causal events. *Cognition*, 88, 1–48.
- Wolff, P., & Song, G. (2003). Models of causation and the semantics of causal verbs. *Cognitive Psychology*, 47, 276–332.

Supporting Information

Additional Supporting Information may be found in the online version of this article on Wiley Online Library:

Table S1. Causal System Category.

Please note: Wiley-Blackwell is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.