

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Identifying causal pathways with and without diagrams

Permalink

<https://escholarship.org/uc/item/8zq0g8qp>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 33(33)

ISSN

1069-7977

Authors

Corter, James
Mason, David
Tversky, Barbara
et al.

Publication Date

2011

Peer reviewed

Identifying causal pathways with and without diagrams

James E. Corter (jec34@columbia.edu)

Teachers College, Columbia University, 525 W. 120th Street
New York, NY 10027 USA

David L. Mason (dml2153@columbia.edu)

Teachers College, Columbia University, 525 W. 120th Street
New York, NY 10027 USA

Barbara Tversky (btversky@columbia.edu)

Teachers College, Columbia University, 525 W. 120th Street
New York, NY 10027 USA

Jeffrey V. Nickerson (jnickerson@stevens.edu)

Stevens Institute of Technology, Castle Point on Hudson
Hoboken, NJ 07030 USA

Abstract

Causal modeling generally involves the construction and use of diagrammatic representations of the causal assumptions, expressed as directed acyclic graphs (DAGs). Do such graphs have cognitive benefits, for example by facilitating user inferences involving the underlying causal models? In two empirical studies, participants were given a set of causal assumptions, then attempted to identify all the causal pathways linking two variables in the model implied by these causal assumptions. Participants who were provided with a path diagram expressing the assumptions were more successful at identifying indirect pathways than those given the assumptions in the form of lists. Furthermore, the spatial orientation of the causal flow in the graphical model (left to right or right to left) had effects on the speed and accuracy with which participants made these inferences.

Keywords: causal models; causal reasoning; causal inference; path models; directed graphs; networks; indirect effects

Introduction

Why are diagrams so much used, and so useful, in learning and reasoning about abstract relationships? Diagrams and language are two ways of externalizing thought to reduce memory load and facilitate inferences (Scaife & Rogers, 1996). Diagrams have the advantage that they can use elements in space and spatial relations to express the elements and relations of thought. Then people's well-developed skills at making spatial judgments and inferences can be applied to abstract judgments and inferences (e.g., Larkin & Simon, 1987; Tversky, 2001). Written language overcomes the fleeting nature of speech, but the form of written language generally bears no direct correspondence to the forms of thought. On the contrary, written language requires people to construct and hold mental models of the forms of thought as well as to use the mental models for reasoning and inference, a double burden that taxes limited working memory capacity. For many reasoning tasks, then, constructing a diagram should both

alleviate limited memory and facilitate inference-making. Abstraction (Schwartz, 1995) and transfer (Novick & Hmelo, 1994) are two types of inferences, ordinarily difficult for people, that may be facilitated by diagrams. It should be noted, however, that diagrams are not always useful for inference and problem-solving; one reason is that learning to construct and use appropriate diagrammatic representations can be difficult (e.g., Ainsworth, 2006; Corter, Nickerson, Tversky, Zahner & Rho, 2008; Zahner & Corter, 2010).

Diagrams are especially appropriate to represent ideas that are inherently or metaphorically spatial, as they readily map elements and relations from some conceptually spatial world to elements and relations on the page. Maps of all kinds, architectural plans, diagrams of the body are examples. Conveying dynamic or invisible properties like change in time, forces, and causes often require the addition of diagrammatic devices like dots, lines, and arrows (Tversky, Zacks, Lee & Heiser, 2000).

Causal modeling is one area where diagrams are conventionally used to represent abstract relationships among entities. Causal modeling involves the use of directed acyclic graphs (DAG) to diagram probabilistic causal relationships. These path diagrams represent variables as nodes and causal relations as directed arrows between pairs of nodes, thus abstracting (and perceptually grounding) the critical information needed for reasoning.

Why do path diagrams play such an integral role in causal modeling? First, some software interfaces for structural equation modeling (SEM) and many systems implementing Bayesian networks (BN) require that the causal model be constructed as a directed graph in a visual programming interface; the resulting diagram is then used to guide the computations for estimation and inference (Greenland, Pearl & Robins, 1999; Lauritzen & Spiegelhalter, 1988; Pearl, 1988). Other researchers have investigated how the structure of the causal network might be inferred from data (e.g., Griffiths & Tenenbaum, 2005; Pearl, 2000; Steyvers,

Tenenbaum, Wagenmakers, & Blum, 2003). But even in traditional approaches to path analysis (e.g., Wright, 1921; Alwin & Hauser, 1975; Bollen, 1989) where the causal network is not directly involved in computation, construction of the path diagram is thought to be an essential step. Causal reasoning in social science domains often involves complex systems of direct and indirect relationships; these concepts may be more easily remembered and understood if these relationships are expressed with diagrams.

However, surprisingly little research has been conducted on what specific cognitive effects and benefits are provided by using path diagrams. McCrudden, Schraw, Lehman, & Poliquin (2007) found that memory for and comprehension of causal relationships from a science text were enhanced by providing a diagram with the text. They also concluded that the benefits of diagrams were greater for more difficult-to-learn causal sequences. Easterday, Alevan, Scheines and Carver (2008) studied the difficulties students have in constructing and interpreting causal diagrams. They found that *providing* students with a diagram for an inference problem about policy options led to best performance on the immediate task, but asking students to *construct* the appropriate diagram led to better transfer performance.

One possible reason that path diagrams help users to reason about complex causal interrelationships is that in a path model, *indirect* influences of variable X on variable Y may involve long causal chains through intervening mediator variables. To use a simple example from educational research, achievement motivation (X) may affect grade-point average (Y) mainly because motivation affects time put into studying (W), which in turn affects Y. Use of the diagrammatic representation makes it easier to find such indirect causal paths involving mediator variables and to correctly interpret how they interact to influence the dependent variable.

Several specific visual aspects of path diagrams may facilitate, or interfere with, the desired inferences. One benefit of the diagram stems from the fact that indirect effects in the causal model (an abstract concept) are represented by *paths* in the diagram, a perceptually basic aspect of the network. Paths in the network are lines; they have the gestalt property of connectedness, meaning that they are perceptually salient and easily understood in relation to our natural abilities to navigate along paths on the two-dimensional surface of the earth. One complication is that both the basic causal assumptions of the model and paths corresponding to indirect effects have directionality. This directionality is indicated in the DAG by arcs or arrows. Arrows are a commonly used device in diagrams, one that is naturally suited for many purposes (Horn, 1998; Kurata & Egenhofer, 2005; Tversky, Zacks, Lee & Heiser, 2000). Arrows are lines, so they connect; but they are asymmetric, indicating an asymmetric relationship such as causation. Students spontaneously interpret arrows in mechanical diagrams as causes, and spontaneously draw arrows in their own visual explanations (Heiser & Tversky,

2006). Here, an arrow indicates an abstraction, that the variable at the tail has a causal influence on the variable at the point.

But previous related research in our lab (Corter et al., 2008; Nickerson, Corter, Tversky, Zahner & Rho, 2008) suggests that several factors might impede efficient use and correct interpretation of path diagrams. These factors are related to the constraints of representation: direct graphs need to be embedded in a two-dimensional page. This fact can of course lead to issues of effective design and use (e.g., how to avoid crossing arcs). More importantly, our previous studies have shown that even though the formal properties of the problem are represented solely by the graph topology, users cannot help being influenced by Euclidean properties of the embedding, such as the distances among nodes (proximity). Humans also show preferences for certain directions in the plane (i.e., left-right and up-down asymmetries), preferences that affect the construction of and processing of external visual representations (Taylor & Tversky, 1992; Tversky, Kugelmass & Winter, 1991; Tversky, 2001).

This brief review suggests a need to replicate and extend the few studies that have investigated cognitive issues surrounding the use and interpretation of path diagrams. The first study described below investigates if use of a path diagram improves users' ability to find causal paths representing indirect causal influences of one variable on another, compared to using text representations of the causal assumptions defining the causal model. We also seek evidence (in the second study) that superficial aspects of the spatial embedding of the path diagram into the plane, in particular left-right directionality, might affect the interpretation and use of the diagram.

Study 1

Do path diagrams improve reasoning about the implications of causal models, specifically the ability to specify all the ways, direct and indirect, in which one variable can causally affect another? Study 1 was designed to provide some initial answers to this question.

Method

In Study 1, we compared how well participants did at identifying all the potential indirect paths between two variables when a path diagram was provided, versus when only a listing of all the assumed direct causal effects was provided.

Participants. Participants were recruited from a crowdsourcing website, Amazon's Mechanical Turk (MT). A total of 172 respondents completed the task. Their mean age was 32.5 ($s = 11.9$), and they were 54% male. Seventy-seven percent of them were native English speakers, and 93% had at least some college education.

Procedure. Participants were randomly assigned to one of three conditions. Two of these conditions presented a set of

causal assumptions in the form of text (in two different formats) and the third presented the same set of assumptions as a path diagram (a directed acyclic graph or DAG).

We first presented participants with instructions that included a worked example, in a format consistent with their assigned condition:

In causal modeling of a social science problem, researchers try to specify all the ways in which variables influence each other. For example, a researcher might assume that variable X affects Y, X affects Z, and Y affects Z. In that case, X has a causal influence on Z in two ways. First, there is a direct effect of X on Z (by assumption). Also, X affects Z indirectly, because X is assumed to affect Y and Y is assumed to affect Z. Thus X has both direct and indirect effects on Z.

This description was followed by one of three displays that depicted the causal assumptions of the preceding problem (*X affects Y, X affects Z, Y affects Z*) as either linear text, text written vertically (as a table), or a path diagram, depending on the condition to which the participant had been assigned. Finally, participants were instructed that the goal was to list all the ways in which variable X could have a causal effect on variable Z (for this example: *X affects Z, and X affects Y which affects Z*). Put in slightly more formal terms, the task was to list all causal pathways (i.e., any direct and all indirect effects) between X and Z.

After reading these instructions and the worked example, participants were then presented with a similar (but more complex) problem to solve. The given information was a set of nine assumptions about pairwise causal relationships among five variables, presented either as horizontal text, vertical text, or as a path diagram in the respective condition (see Figure 1). Participants were asked to “Please write all the ways that variable H could influence variable R.”

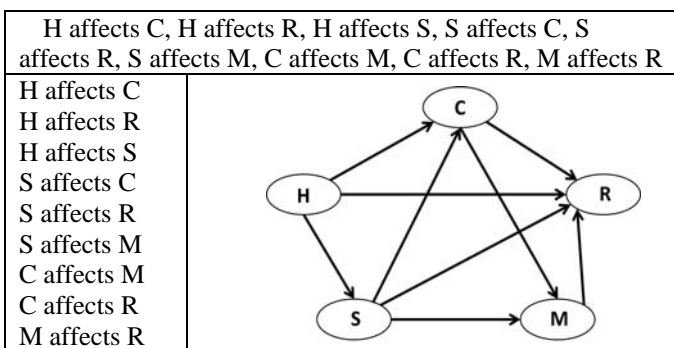


Figure 1: Presentation of a set of causal assumptions in one of three formats: as horizontal text, as vertical text, or as a diagram.

Results

As shown in Table 1, the number of correct responses was highest for the diagram condition (50%), followed by horizontal text (43%), then vertical text (24%). As hypothesized, the difference between the proportions correct

for the diagram condition versus the two text conditions was significant in a log-linear analysis ($z=2.916, p=.004$), as was the difference between vertical and horizontal text ($z = 2.165, p=.030$). As Table 1 shows, incorrect paths were only rarely given; most incorrect answers were due to omitting paths.

Table 1: Study 1 performance (mean accuracy and total work time), by condition.

Condition	N	%corr.	correct paths	incorrect paths	Time (sec)
Horizontal text	53	43	5.9	.06	509.9
Vertical text	59	24	5.3	.08	517.7
Diagram	60	50	6.0	.05	380.9

Table 1 also shows that time required to do the task, as measured by the web-based task administration software, differed among the three conditions, with participants in the diagram condition completing the task marginally faster than participants in the text conditions, $F(1,169) = 3.430, p=.066$.

Not surprisingly, the probability that a participant omitted an indirect effect in their answer tended to increase with the length of the corresponding causal path. The percentage correct for the path of length 2 (the direct effect, HR) was 98%; for paths of length 3 the figure was 92%; for paths of length 4 it was 78%, and for HSCMR, the only path of length 5, it was 54%. The average benefit of using a diagram over the text conditions increased with path length: the advantage in accuracy for paths of length 2, 3, 4 and 5 was 2%, 5%, 5%, and 17% respectively.

In response to an explicit post-task question, 27% of participants in the text conditions reported constructing their own diagrams “offline”, on scratch paper, in the process of answering the paths question, while only 2% reported doing so in the diagram condition. Thus, the observed advantage of the diagram condition over the text conditions is probably underestimated. In the horizontal text condition, 23% reported making a list or table summarizing the assumed causal relationships, versus only 14% in the vertical text condition and 13% in the diagram condition. Consequently, the advantage observed here for horizontal text over vertical text may also be affected by user-generated external representations.

Discussion

As expected, participants in the diagram condition indeed showed higher accuracy in identifying all direct and indirect causal pathways between the two target variables. An unanticipated finding was the large (and significant) advantage of horizontal text over vertical text. However, note that the text presented to participants in this study had a high degree of organization: causal links (pairs) were

presented in an order consistent with the implied causal chains, so that the causal chains do have a kind of indirect visual representation in the horizontal text condition. If the lists of causal assumptions had not been organized in this particular chain-consistent order, the causal chains would likely have been less salient, and the advantage of diagram over text would likely have been even higher.

Study 2

Study 2 was an attempt to replicate the main findings of Study 1 and assess the generalizability of the diagram advantage by using a new set of causal assumptions corresponding to a new diagram structure. Also, the effect of manipulating one superficial aspect of the diagram was investigated, in this case whether the causal flow in the diagram was generally from left to right (the conventional orientation observed in many English texts and journal articles on causal modeling), or from right to left. Our previous work (Corter et al., 2008; Nickerson et al., 2008) has shown cognitive effects of supposedly superficial aspects of how network diagrams are embedded on the page, including preferences for top-down and left-right processing of diagrams (see also Taylor & Tversky, 1992; Tversky et al., 1991).

Method

The methods of Study 2 were essentially the same as for Study 1, except for use of a new causal structure, and the addition of a variant diagram with right-to-left causal flow.

Participants. Participants were recruited from a crowd-sourcing website. After eliminating participants who had participated in Study 1, and those who failed to follow instructions for the Study 2 task, we were left with N=212 participants. Their mean age was 31.1 (s = 11.1), and they were 54% male. Seventy-six percent of them were native English speakers, and 89% had at least some college education.

Procedure. Task instructions for Study 2 used the same worked example as did the instructions for Study 1. The only substantive change to the procedure (besides use of a different causal structure) involved adding a second diagram condition, in which causal flow in the diagram proceeded from the right side of the diagram to the left, rather than left to right, as is usual practice. Thus, the tested conditions were: horizontal text, vertical text, diagram l-r, and diagram r-l (see Figure 2). Participants were asked to “Please write all the ways that variable C could influence variable S.”

Results

The proportion of correct answers again differed among conditions (Table 2), with the two diagram conditions showing higher performance than the two text conditions. This advantage was confirmed in a log-linear analysis, $z = 2.290$, $p = .022$. Surprisingly, the right-to-left diagram resulted in higher accuracy (62%) than the left-to-right

diagram (52%), though this difference was not significant, $z = 1.014$, $p = .312$. However, the total work times reported in Table 2 reveal that participants worked much more slowly with the right-to-left version. Accuracy did not differ between the two text conditions, $z = 0.297$, $p = .764$.

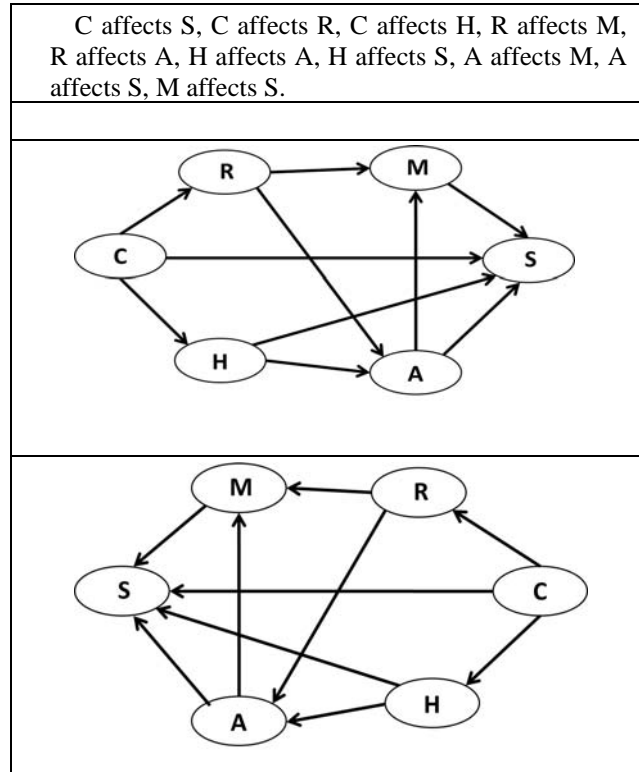


Figure 2: Presentation of a set of causal assumptions (Study 2) as horizontal text, as a diagram with left-to-right causal flow, and with right-to-left causal flow (vertical text condition not shown).

Table 2: Study 2 performance (mean accuracy and total study work time), by condition.

Condition	N	% correct	# corr. Paths	Time (sec)
Horiz. text	46	39	5.78	526.0
Vertical text	45	42	6.02	617.0
Diagram l-r	69	52	6.33	355.9
Diagram r-l	52	62	6.21	620.8

As in Study 1, the probability that participants omitted an indirect effect in their answers increased with the length of the causal path. The overall percentage correct for the direct effect, (path SC) was 98%; for path (CHS) it was 96%; for paths of length 4 it was 88%, and for length 5 paths, it was 74%. The advantage due to using a diagram tended to increase with path length, though the advantage for paths of length 4 (9%) exceeded that for paths of length 5 (4%).

In Study 2 31% of participants in the text conditions reported constructing their own diagrams offline in the process of answering the paths question, while only 6% reported doing so in the diagram conditions. This is further evidence, albeit indirect, that diagrams are useful to participants trying to answer the inference question, and suggests that the higher accuracy observed here for the diagram conditions compared to the text conditions may be an underestimate of the true effect of using diagrams. Also, in the text conditions 14% of participants reported making their own list or table of the assumed causal relationships, versus 10% in the diagram conditions.

Discussion

Again, the results show that participants did a better job of identifying all the indirect causal paths from one variable to another when they were provided with a diagram representing the underlying causal assumptions. Furthermore, many participants who were not provided with a diagram reported constructing one on their own, presumably to aid themselves in the task. Participants were slowed in their work when the causal diagram presented the causal flow from right-to-left, an unusual orientation; however, this manipulation seemed to actually improve accuracy. This effect, should it prove replicable, might be due to the unfamiliar orientation triggering more careful processing (e.g., Bjork, 1994; Alter, Epley, Oppenheimer & Eyre, 2007).

General Discussion

This research adds to the body of evidence that diagrams are useful external aids to reasoning. In both studies reported here, providing participants with a path diagram improved their accuracy in finding all direct and indirect effects of one variable on another, a task that is equivalent to specifying all the causal paths between those variables.

It is worth re-emphasizing that the advantages in inference accuracy found here for diagrams are likely to be underestimates of the true benefits, for several reasons. First, many participants in the text conditions reported constructing their own diagram “offline” in answering the inference question, even though they were not asked to do so. Also, note that the text versions of the problems were highly organized in a way that should promote the finding of causal paths. Because the models examined here were recursive, and the given direct causal relationships between pairs of variables were presented in a lexicographic order based on (cause, effect), each causal path corresponding to an indirect effect (i.e., the component paths in the correct answer) could be constructed with a single “pass” through the list of causal relationships. Because this degree of organization in the text versions of the problems is both optimal and artificial (or at least a special case), organizing the lists in any other way could be expected to lower performance for the text conditions, increasing the measured benefit of diagrams.

Another noteworthy finding, from Study 2, is that use of the path diagram seemed to be affected by a superficial aspect of how the causal diagram was embedded into the plane – specifically, by whether the causal flow was depicted as generally from left-to-right or right-to-left. In previous work involving how people reason using network diagrams (Corter et al., 2008; Nickerson et al., 2008), we have found that people use and interpret such supposedly superficial aspects of how the formal diagram is embedded on the two-dimensional page (c.f. Landy & Goldstone, 2007). The typical left-to-right reading of diagrams often displayed by native speakers of Western languages (Taylor & Tversky, 1992; Tversky, Kugelmass & Winter, 1991) is a predilection that diagram designers may take into account. Thus, these findings add to the growing evidence that diagram designers and users are affected by and take advantage of the affordances of the page.

Acknowledgments

This work has been supported by awards IIS-0725223, HHC-0905417, and IIS-0855995 from the National Science Foundation, and by the Stanford Regional Visualization and Analysis Center.

References

- Ainsworth, S. (2006). DeFT: A conceptual framework for considering learning with multiple representations. *Learning and Instruction, 16*, 183-198.
- Alter, A. L., Oppenheimer, D. M., Epley, N., & Eyre, R. N. (2007). Overcoming intuition: Metacognitive difficulty activates analytic reasoning. *Journal of Experimental Psychology: General, 136*(4), 569–576.
- Alwin, D. F., & Hauser, R. M. (1975). The decomposition of effects in path analysis. *American Sociological Review, 40*, 37-47.
- Bjork, R.A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe and A. Shimamura (Eds.) *Metacognition: Knowing about knowing*. (pp.185-205). Cambridge, MA: MIT Press
- Bollen, K. A. (1989). *Structural Equations with Latent Variables*. New York: Wiley.
- Corter, J. E., Nickerson, J.V., Tversky, B., Zahner, D., and Rho, Y. (2008). Using diagrams to design information systems. In B. C. Love, K. McRae, and V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, (pp. 2259–2264). Austin, TX: Cognitive Science Society.
- Easterday, M. W., Aleven, V., & Scheines, R. (2007). ‘Tis better to construct than to receive? The effects of diagram tools on causal reasoning. *Proceedings of AIED*, pp. 93-100, IOS Press, Amsterdam.
- Greenland, S., Pearl, J., & Robins, J. M. (1999). Causal diagrams for epidemiologic research. *Epidemiology, 10*(1), 37-48.

- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*, 334–384.
- Heiser, J. & Tversky, B. (2006). Arrows in comprehending and producing mechanical diagrams. *Cognitive Science*, *30*, 581-590.
- Landy, D., & Goldstone, R. L. (2007). How abstract is symbolic thought? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(4), 720–733.
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, *11*, 65-99.
- Lauritzen, S., and Spiegelhalter, D. (1988). Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society*, *50*, 157–224.
- McCrudden, M. T., Gregory Schraw, G., Lehman, S., & Poliquin, A. (2007). The effect of causal diagrams on text learning. *Contemporary Educational Psychology*, *32*, 367–388.
- Nickerson, J. V., Corter, J. E., Tversky, B., Zahner, D., and Rho, Y. (2008). The spatial nature of thought: Understanding information systems design through diagrams. In Boland, R., Limayem, M., and Pentland B., (Eds.), *Proceedings of the 29th International Conference on Information Systems*.
- Novick, L. R., & Hmelo, C. E. (1994). Transferring symbolic representations across nonisomorphic problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(6), 1296-1321.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo CA: Morgan Kaufmann.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge: Cambridge University Press.
- Scaife, M., & Rogers, Y. (1996). External cognition: How do graphical representations work? *International Journal of Human-Computer Studies*, *45*, 185-213.
- Schwartz, D. L. (1995). The emergence of abstract representations in dyad problem solving. *The Journal of the Learning Sciences*, *4*(3), 321-354.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–489.
- Taylor, H.A., & Tversky, B. (1992). Descriptions and depictions of environments. *Memory & Cognition*, *20*, 483-496.
- Tversky, B., Kugelmass, S., & Winter, A. (1991) Cross-cultural and developmental trends in graphic productions. *Cognitive Psychology*, *23*, 515-557.
- Tversky, B. (2001). Spatial schemas in depictions. In M. Gattis (Ed.), *Spatial schemas and abstract thought*, Pp. 79–111. Cambridge: MIT Press.
- Tversky, B. (2005). Prolegomenon to scientific visualizations. In J. K. Gilbert (Ed.), *Visualization in science education*, Series: *Models and modeling in science education*, Dordrecht, Kluwer, 1, 29-42.
- Wright, S. (1921). Correlation and causation. *Journal of Agricultural Research*, *20*, 557-585.
- Zahner, D., & Corter, J. E. (2010). The process of probability problem solving: Use of external visual representations. *Mathematical Thinking and Learning*, *12*(2), 177 - 204.