

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A Vector Model of Causal Meaning

Permalink

<https://escholarship.org/uc/item/1qm8q0rt>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 24(24)

ISSN

1069-7977

Authors

Wolff, Phillip
Zettergren, Matthew

Publication Date

2002

Peer reviewed

A Vector Model of Causal Meaning

Phillip Wolff (pwolff@memphis.edu)

Department of Psychology, 202 Psychology Bldg.
Memphis, TN 38152 USA

Matthew Zettergren (mzttgrn@memphis.edu)

Department of Electrical and Computer Engineering, Engineering Bldg.
Memphis, TN 38152 USA

Abstract

This paper proposes a new model of causal meaning, the Vector Model, which formalizes a model of causation based on Talmy's notions of force dynamics (Wolff, Song, & Driscoll, 2002). In the Vector Model, the concepts of CAUSE, ENABLE and PREVENT are distinguished from one another in terms of force vectors, their resultant and the relationship of each force vector to a target vector. The predictions of the model were tested in two experiments in which participants saw realistic 3D-animations of an inflatable boat moving through a pool of water. The boat's movements were completely determined by the force vectors entered into a physics simulator. Participants' linguistic descriptions of the animations were closely matched by those predicted by the model given the same force vectors as those used to produce the animations. Our model may have implications for the semantics of causal verbs as well as the perception of causal events.

Introduction

This research investigates people's notions of causation as reflected in their use of causal verbs. We approach this problem by formulating a model of causal meaning that defines causal concepts in terms of relationships between force vectors, their resultant and a target position vector.

We begin by noting two key problems for models of causal meaning. First, such models must be able to distinguish the concept of CAUSE from the concept of ENABLE. We say, for example, the wave (and not the keel) *caused* the sailboat to rock, while the keel (and not the wave) *enabled* the sailboat to rock. The precise way in which these two notions differ has been difficult to specify. Contributing to this difficulty is the fact that the two concepts cannot be distinguished in terms of necessity or sufficiency (Cheng & Novick, 1991; Goldvarg & Johnson-Laird, 2001). In the above example, neither the wave nor the keel alone is sufficient, but both may be necessary for the boat's rocking to occur. Several solutions to this challenge have been proposed, but most have not escaped criticism (see Cheng & Novick, 1991; Goldvarg & Johnson-Laird, 2001; Wolff, Song, & Driscoll, 2002).

A second key problem for models of causal meaning concerns how the concept of CAUSE is represented in expressions that refer to specific instances of causation. Many models of causation define causation in terms of probabilities (e.g., Cheng, 1997; Cheng & Novick, 1991; Glymour, 2001). Such models are well suited for explaining the meaning of generic statements of causation, that is, statements about what is typically the case in multiple occurrences of a particular event, as in *Heavy snowmelt causes rivers to flood*. What these theories do not handle well are expressions that refer to a single instance of causation, as in *The heavy snowmelt caused the Colorado to flood*. Sentences describing single instances express what is definitely true of a particular event, not what is typically true of many. Moreover, such sentences are incompatible with the non-occurrence of the result (e.g., flooding), but if causation is inherently probabilistic, such non-occurrences cannot be strictly ruled out (Goldvarg & Johnson-Laird, 2001).

In some theories of causation, the concept of CAUSE is defined in such a way that it can be used in descriptions of singular causation. For example, according to Michotte (1963), causation is inferred from the perception of a transfer of motion from one ball to another—an “ampliation of motion” (p. 143). A related proposal is that CAUSE is inherently based on the idea of force and that the occurrence of CAUSE involves a mechanism by which this force is transmitted (Ahn & Kalish, 2000; Shultz, 1982). While these theories specify properties that could be predicated of a single event (and are highly related to the proposal we make in this paper), they do not provide us with a clear solution to the first problem of causal meaning: how the notion of CAUSE might be distinguished from the notion of ENABLE.¹ Both CAUSE and ENABLE presumably involve the transference of force.

In this paper, we propose a model of causal meaning that addresses these two problems. This model represents a formalization of the Force Dynamic Model described in Wolff, Song and Driscoll (2002; also Wolff & Song, 2001). In the next section, we describe

¹ Counterfactual theories of causation face related problems (see Spellman & Mandel, 1999)

the Force Dynamic Model as well as some of the empirical evidence in support of it. We then turn to a description of its formalization.

The Force Dynamic Model of Causation

A theory of force dynamics was first proposed by Talmy (1988), and has been elaborated by several other researchers (Jackendoff, 1991; Kemmer & Verhagen, 1994; Pinker, 1989; Robertson & Glenberg, 1998; Siskind, 2000; Verhagen & Kemmer, 1997). From a force dynamic perspective, the concept of CAUSE is one member of a family of concepts that include the concepts of ENABLE and PREVENT, among others. With each of these concepts, there are two key players: an affector and a patient.² Differences among the concepts are captured in terms of various patterns of tendency, relative strength, rest, and motion.

The Force Dynamic Model specified in Wolff, Song and Driscoll (2002) combines two of Talmy's (1988) core dimensions (Tendency & Result) with a dimension suggested by Jackendoff (1991).³

Table 1: The Force Dynamic Model's representations of CAUSE, ENABLE, & PREVENT

	Patient Tendency for Result	Affector-Patient Opposition	Occurrence of Result
CAUSE	N	Y	Y
ENABLE	Y	N	Y
PREVENT	Y	Y	N

As shown in Table 1, this model specifies that the concepts of CAUSE, ENABLE, and PREVENT can be captured in terms of 1) the tendency of the patient for the result, 2) the presence of opposition between the affector and patient, and 3) the occurrence of the result. In causing situations (see 1a), for example, the tendency of the patient, the boat, is not for the result, heeling. But because the tendency is opposed by the affector, the result, i.e., heeling, occurs.

- (1) a. The blast caused the boat to heel.
- b. Vitamin B enables the body to digest food.
- c. The rain prevented the tar from bonding.

In enabling situations, as in (1b), the tendency of the patient, the body, is for the result, to digest food. This tendency is not opposed by vitamin B. Rather, vitamin B assists in the realization of this tendency, which leads to the occurrence of a result. In situations involving preventing, as in (1c), the tendency of the patient, the tar, is towards the occurrence of the result, bonding, but

this tendency is opposed and blocked by the affector, and as a consequence, the result does not occur.

Evidence in support of the Force Dynamic Model

As indicated in Table 1, the Force Dynamic Model predicts that each concept shares one feature in common with each other concept: ENABLE and PREVENT both involve patients with a tendency for the result; CAUSE and PREVENT both involve opposition; and CAUSE and ENABLE both lead to results. The model implies, then, that the three concepts should be equally similar in meaning. Therefore, if we were to plot these concepts in a similarity space in terms of the verbs that encode them, they should reside roughly equally distant from one another. In fact, this is exactly what we found when we asked people to sort 48 sentences from the British National Corpus that contained 23 periphrastic causative verbs (i.e. verbs that pattern syntactically and semantically like the verb *cause*, e.g., *make*, *enable* and *prevent*) and submitted their sorts to a multidimensional scaling program⁴ (Wolff et al., 2002).

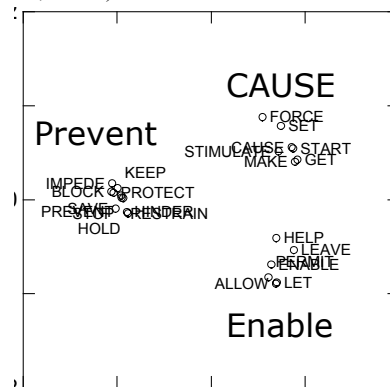


Figure 1: MDS solution of periphrastic causative verbs

As Figure 1 shows, the periphrastic causative verbs in English fall into three categories: a CAUSE category that includes the verbs *cause*, *force*, *get*, *make*, *set* and *stimulate*, an ENABLE category that includes the verbs *allow*, *enable*, *help*, *leave*, *let*, and *permit*, and a PREVENT category that includes the verbs *block*, *hinder*, *hold*, *impede*, *keep*, *prevent*, *protect*, *restrain*, *stop*. Importantly, the clusters associated with these three concepts reside roughly equally distant from one another, just as predicted by the Force Dynamic Model. We have replicated these results for specific and generic statements of causation. These results, along with several rating studies, lead us to believe that the Force Dynamic Model captures the primary semantic dimensions underlying the periphrastic causative verbs, and the verb *cause* in particular.

² We use the more familiar terms affector and patient instead of antagonist and agonist as originally used in Talmy (1988).

³ In Talmy (1988) nearly all interactions involve opposition while in Jackendoff (1991) this parameter is allowed to vary.

⁴ Multidimensional scaling is a procedure that locates items in space so that their distances in that space reflect as closely as possible their measured inter-item (dis)similarities.

The Vector Model of Causation

In the Vector Model, the notions of tendency, opposition (here, concordance), and result are represented as force vectors, their resultant and the relationship of each force vector to a target position vector. The model is described below for physical interactions in which the patient has no initial velocity. However, it is assumed that it could be extended to situations in which the patient does have an initial velocity (and, hence, momentum). It is also assumed that the model could be extended to cover non-physical kinds of causation (e.g., social, psychological).

In our description, all vectors are typed in boldface font; $\mathbf{P} \cdot \mathbf{T}$ denotes the dot product of the vectors \mathbf{P} and \mathbf{T} ; $\|\mathbf{P}\|$ denotes the magnitude of \mathbf{P} .

In the case of physical causation, \mathbf{A} represents a vector that specifies the force exerted on the patient by the affector; \mathbf{P} , any force produced by the patient to move itself, or in the absence of such a force, its weight (e.g., force pulling it toward the earth) and/or resistance to motion due to frictional forces; \mathbf{O} , the vector representing the summation of the remaining *other* forces acting on the patient⁵ and \mathbf{R} , the resultant force acting on the patient based on the vector addition of \mathbf{A} , \mathbf{P} and \mathbf{O} . An example configuration is shown in Figure 2.




Figure 2: Forces associated with the affector, \mathbf{A} , patient, \mathbf{P} , and other forces, \mathbf{O} , combine to produce a resultant force, \mathbf{R} , in the direction of a target.

The target's location is specified in terms of a position vector, \mathbf{T} . When the target and patient are points, \mathbf{T} simply begins at the patient and ends at the target, as shown in Figure 3.




Figure 3: The target's location is specified by a position vector \mathbf{T} .

In the more general case in which the target is represented by an area, the target's location would be specified by a set of real 1- or 2-dimensional position

⁵ The contribution of other forces, \mathbf{O} , might include forces whose entities could serve as affectors or patients in other interactions, as well as forces that might be used to distinguish between periphrastic causative verbs within a subcategory (e.g., *help* vs. *enable* vs. *allow* vs. *let*).

vectors, such that every vector from the patient's position to a point that could be considered a part of the target would be an element of that set.⁶

For this particular version of the Vector Model, we assume that all of the forces are constant with respect to time and space (i.e., $\partial/\partial t[\mathbf{Z}(x,y,t)] = \partial/\partial x[\mathbf{Z}(x,y,t)] = \partial/\partial y[\mathbf{Z}(x,y,t)] = \mathbf{0}$ where \mathbf{Z} is any force in this model), the patient has no initial velocity, and $\|\mathbf{P}\|$ and $\|\mathbf{T}\| > 0$.

The main dimensions of the Vector Model are defined in Table 2 for patients and targets that can be construed as points, and where $\|\mathbf{A}\| > 0$.

Table 2: Dimensions underlying the Vector Model

Dimension	Formal Definition
<i>Tendency</i> (of patient for the target)	Angle between \mathbf{P} and $\mathbf{T} = 0^\circ$
<i>Concordance</i> (of affector & patient)	Angle between \mathbf{A} and $\mathbf{P} = 0^\circ$
<i>Result</i>	Angle between \mathbf{R} and $\mathbf{T} = 0^\circ$

Rationale for the definitions *Tendency* - If the patient has a tendency for the target, then the direction of its force vector will coincide with the direction of the position vector \mathbf{T} . Thus, the angle between the vector \mathbf{P} and \mathbf{T} will be 0° . A test for this possibility can be stated with respect to the dot product of \mathbf{P} and \mathbf{T} . Specifically, when the patient has a tendency for the target, $\mathbf{P} \cdot \mathbf{T} = \|\mathbf{P}\| \cdot \|\mathbf{T}\|$ ⁷, and when it does not, $\mathbf{P} \cdot \mathbf{T} < \|\mathbf{P}\| \cdot \|\mathbf{T}\|$.

Concordance - Concordance concerns the similarity of the force vectors associated with the affector and the patient. If the affector and patient exert forces (on the patient) in the same direction, then they are considered to be in concordance. In a similar fashion to tendency, concordance can be defined with respect to the dot product, but this time between \mathbf{P} and \mathbf{A} . Specifically, when the affector and patient are in concordance, $\mathbf{P} \cdot \mathbf{A} = \|\mathbf{P}\| \cdot \|\mathbf{A}\|$, and when they are not, $\mathbf{P} \cdot \mathbf{A} < \|\mathbf{P}\| \cdot \|\mathbf{A}\|$ ⁸.

Result - As with tendency and concordance, occurrence of a result can be defined in terms of the similarity between two vectors, but this time between \mathbf{R} and \mathbf{T} . When the angle between \mathbf{R} and \mathbf{T} is 0° , the result will occur, assuming all of the forces acting on

⁶ In the more general case in which the target is other than a point, we expect the definition of concordance must be changed to include a certain level of angular tolerance that would be based, in part, upon the relative size of the target and its proximity to the patient.

⁷ By definition of the dot product, $\mathbf{P} \cdot \mathbf{T} = \|\mathbf{P}\| \cdot \|\mathbf{T}\| \cdot \cos(\theta)$, where θ is the angle between vectors \mathbf{P} and \mathbf{T} . In the case where θ is 0° , the equation becomes $\mathbf{P} \cdot \mathbf{T} = \|\mathbf{P}\| \cdot \|\mathbf{T}\| \cdot \cos(0^\circ)$, which reduces to $\mathbf{P} \cdot \mathbf{T} = \|\mathbf{P}\| \cdot \|\mathbf{T}\|$.

⁸ When concordance is defined in terms of the dot product, it allows for a special type of concordance in which $\|\mathbf{A}\| = 0$. When $\|\mathbf{A}\| = 0$, the equality $\mathbf{P} \cdot \mathbf{A} = \|\mathbf{P}\| \cdot \|\mathbf{A}\|$ would hold, which may be representative of the kinds of situations referred to by the verbs *let*, *allow*, and *permit*.

the patient are constant with respect to time and space, as specified formally above. In terms of the dot product, the result will occur if $\mathbf{R} \cdot \mathbf{T} = \|\mathbf{R}\| \|\mathbf{T}\|$ and will not occur if $\mathbf{R} \cdot \mathbf{T} < \|\mathbf{R}\| \|\mathbf{T}\|$.

As with the Force Dynamic Model, CAUSE, ENABLE, and PREVENT are defined with respect to values along three dimensions, specified in Table 3, and share one feature with each other concept. Thus, both models predict that the three concepts should be equally similar to one another.

Table 3: The Vector Model’s representations of CAUSE, ENABLE, & PREVENT

	Tendency of Patient for Target	Concordance of Affector & Patient	Result
CAUSE	N	N	Y
ENABLE	Y	Y	Y
PREVENT	Y	N	N

Testing the Vector Model of Causation

Beyond similarity, the Vector Model makes predictions about the vector configurations underlying verbs of causation. These predictions were tested in two experiments. Participants viewed 3D animations of an inflatable boat, the patient, moving across a shallow pool in relationship to a half-submerged cone, the target (see Figure 4). Each animation had two main parts. In the first, the boat moved from the side of the pool to the center. This part was included to establish the boat’s tendency. In the second part, a bank of fans (i.e., the affector) started blowing. Thus, in the second part of every animation, the force produced by the boat itself was combined with the force exerted on the boat by the fans to give rise to a resultant force that determined the boat’s direction and speed.

After watching an animation, participants chose among several possible linguistic descriptions. We predicted, per the Vector Model (and its computer implementation), that participants would choose a description containing the verb *cause* when the boat started moving away from the cone (Tendency = N), but was moved to the cone (Result = Y) by the fans blowing in a direction different from the direction of the boat (Concordance = N). We predicted that participants would choose a description containing the verb *help* (a type of ENABLE verb, see Figure 1), when the boat moved towards the cone (Tendency = Y) and ultimately reached it (Result = Y) when the fans blew in the same direction as the boat’s direction of motion (Concordance = Y). We also predicted that participants would choose a description containing the verb *prevent* when the boat started towards the cone (Tendency = Y) but did not hit it (Result = N) because the fans blew it back or away from the cone (Concordance = N). Finally, we predicted that when none of the above

configurations were instantiated, participants would choose “none of the above.” These predictions were tested for one- and two-dimensional interactions in Experiments 1 and 2, respectively.



Figure 4: Sample frame from an animation used in Experiment 2 that instantiated a “cause” interaction

Experiment 1

Method

Participants The participants were 18 University of Memphis undergraduates.

Materials Eight 3D animations were made from an animation package called Discreet 3ds max 4. The direction and speed of the boat was calculated by a physics simulator called Havok Reactor. In each animation the boat was initially located four boat-lengths away from the center of the pool. In the first half of the animation, the boat moved towards the center, ostensibly under its own power. Once the boat reached the center, the fans started blowing. The animation ended when the boat hit the cone or neared the side of the pool (~4 seconds total).

The top of Table 4 shows the direction and relative magnitudes of the force vectors associated with the affector and patient that were entered into the physics simulator. The affector, **A**, and patient, **P**, vectors were either in the direction of the target or in the opposite direction. The magnitude of the other forces vector, **O**, was set to 0. In half of the interactions, the affector vector was 1.7 times stronger than the patient (configurations 1-4), while in the remaining interactions the strengths were reversed (configurations 5-8).

Procedure The animations were presented in random order on Windows-based computers. After each animation, participants chose a sentence that best described the occurrence. All of the sentences were the same (“The fans ____ the boat to [from] hit[ting] the cone”) except for the verb, which was either *caused*, *helped* or *prevented*. Another option was “none of the above.” Participants indicated their answers by clicking a radio button next to their choice.

Table 4. The vectors configurations used in Experiment 1, along with associated predictions and results

Configuration	1	2	3	4	5	6	7	8
Affector (\Rightarrow)								
Patient (\rightarrow)								
Target (T)								
Predictions	<i>Help</i>	<i>Cause</i>	<i>Prevent</i>	No verb	<i>Help</i>	No verb	No verb	No verb
Results								
<i>Cause</i>	11%	94%	-	-	6%	6%	-	-
<i>Help</i>	89%	6%	-	-	94%	-	11%	-
<i>Prevent</i>	-	-	100%	-	-	-	6%	6%
No verb	-	-	-	100%	-	94%	83%	94%

Results and Discussion

The predictions of the Vector Model were fully borne out by the results. The bottom of Table 4 shows the percentage of times people chose each of the four possible options for each of the vector configurations. Participants chose *cause*—as opposed to the other possible options—for the animation in which the boat first moved away from the cone but was later pushed back against it by the fans (configuration 2), a N-N-Y type of occurrence in terms of tendency, concordance and result (see Table 3), $\chi^2(3, N=18) = 62, p < .001$. Participants chose *help* when the direction of the boat and the fans was the same (1, 5), a Y-Y-Y type of occurrence, $\chi^2(3, N=18) = 116, p < .001$. Participants chose *prevent* when the boat moved towards the cone but was then kept from hitting it by the fans (3), a Y-N-N occurrence, $\chi^2(3, N=18) = 72, p < .001$. Finally, participants chose “none of the above” when the vector configurations did not map onto any one of the three main kinds of configurations, $\chi^2(3, N=18) = 237, p < .001$. Importantly, participants did not choose *prevent* whenever the boat missed the cone (4, 6, 8). Instead, *prevent* was restricted to those situations in which the boat had an initial tendency for the target (3). Likewise, participants did not choose *cause* or *enable* when the boat simply hit the cone (7), but only when the vector configurations matched those defined by the model. Thus, the Vector Model is capable of not only specifying distinct types of causal concepts, but also distinguishing between causation and non-causation.

The results strongly support the Vector Model, but only in the case of interactions occurring within a single dimension. In Experiment 2 we examine the ability of the model to handle two-dimensional interactions.

Experiment 2

Method

Participants. The participants were 18 University of Memphis undergraduates.

Materials. Ten 3D animations were made in the same way as in Experiment 1 except that the affector and patient force vectors were oriented in several directions other than directly towards or away from the target, and the magnitudes of the affector and patient vectors were always the same. The ten vector combinations at the top of Table 5 depict five combinations in which the patient vector is oriented away from the target by 45° (1-5) and five combinations in which the patient vector is oriented towards the target (6-10). The affector vector was oriented from 180° to 360° at 45° intervals.

Procedure The procedure was as in Experiment 1.

Results

The predictions of the Vector Model were supported once again. The bottom of Table 5 shows the percentage of times people chose each of the four possible options for each of the vector configurations. Participants chose *cause* for the animation in which the

Table 5. The vectors configurations used in Experiment 2, along with associated predictions and results

Configuration	1	2	3	4	5	6	7	8	9	10
Affector (\Rightarrow)										
Patient (\rightarrow)										
Target (T)										
Predictions	No verb	<i>Cause</i>	No verb	No verb	No verb	<i>Help</i>	<i>Prevent</i>	<i>Prevent</i>	<i>Prevent</i>	<i>Prevent</i>
Results										
<i>Cause</i>	-	89%	-	-	-	11%	-	-	-	-
<i>Help</i>	-	11%	-	-	-	83%	-	-	-	-
<i>Prevent</i>	-	-	17%	-	11%	-	94%	94%	89%	89%
No verb	100%	-	83%	100%	89%	6%	6%	6%	11%	11%

boat was not headed for the cone but hit it because of the fans (2, a N-N-Y occurrence), $\chi^2(3, N=18) = 53, p < .001$. Participants chose *help* when the boat was headed for the cone and then was assisted in hitting it by the fans (6, a Y-Y-Y occurrence), $\chi^2(3, N=18) = 44, p < .001$. Participants chose *prevent* when the boat was initially headed toward the cone but was later blown away from it (7, 8, 9, 10, a Y-N-N occurrence), $\chi^2(3, N=18) = 229, p < .001$. Finally, participants chose “none of the above” when the vector configurations did not map onto any one of the three main kinds of configurations (1, 3, 4, 5), $\chi^2(3, N=18) = 238, p < .001$.

Conclusions

In this research we proposed a new model of causal meaning. We also provided empirical support for this model by showing that people’s linguistic descriptions of animations are well accounted for by the model and its computer implementation given the same force vectors as those used to produce the animations.

According to the Vector Model, each kind of causal relation is associated with a range of spatial geometries in addition to a particular temporal organization. As a consequence, the model is able to handle causal relations that are highly problematic for probabilistic models, in particular, those in which the cause and effect occur simultaneously (*The sun’s gravity causes the earth to revolve around the it*). In such situations, it is difficult to count the causing and resulting events for the purposes of calculating probabilities. In contrast, for the Vector Model, such situations are not problematic since they give rise to readily identifiable vector configurations.

The model provides a new explanation for why billiard-ball events, like the ones studied by Michotte (1963), are construed as causal. Traditionally, this was explained in terms of the spatial-temporal contiguity of the causing and resulting events. Clearly, spatial-temporal contiguity is important: without it, there can be no interaction of (contact-type) forces. But spatial-temporal contiguity is not particular to causal interactions alone. According to the Vector Model, what leads people to describe billiard-ball events as causal is that the patient resists moving (Tendency=N), the affector opposes this tendency (Concordance=N), and the patient ends up moving (Result=Y).

In sum, the Vector Model is able to address several important problems in the causation literature in addition to the two problems discussed in the introduction: the distinction between CAUSE and ENABLE and the expression of singular causation. It also takes us a step closer towards understanding how physical interactions may be construed for the purposes of language.

Acknowledgments

This work was supported by grants award to the first author from The University of Memphis Faculty Research Fund and ONR (N00014-01-1-0917). The conclusions of this research do not necessarily reflect those held by these funding sources. We give special thanks to Tanya Vassilieva, Bianca Klettke, and Derek Wong for their help in the research.

References

- Ahn, W., & 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.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psych. Review*, *104*, 367-405.
- Cheng, P. W. & Novick, L. R. (1991). Causes versus enabling conditions. *Cognition*, *40*, 83-120.
- Goldvarg, E., & Johnson-Laird, P. (2001). Naive Causality: a mental model theory of causal meaning and reasoning. *Cognitive Science*, *25*, 565-610.
- Jackendoff, R. (1991). *Semantic Structures*. Cambridge, MA: The MIT Press.
- Kemmer, S. & Verhagen, A. (1994). The grammar of causatives and the conceptual structure of events. *Cognitive Linguistics*, *5*, 115-56.
- Michotte, A. (1963). *The perception of causality*. London: Methuen. (Originally published 1946.)
- Glymour, C. (2001). *The mind’s arrows*. MIT Press.
- Pinker, S. (1989). *Learnability and Cognition*. Cambridge, MA: MIT Press.
- Robertson, D. A., & Glenberg, A. M. (1998). Force dynamics in language and cognition. *Proceedings of the 20th annual meeting of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Shultz, T. R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development*, *47*, 1-51.
- Siskind, J. M. (2000). Visual event classification via force dynamics. *Proceedings AAAI-2000*, 149-55.
- Spellman, B. A., & Mandel, D. R. (1999). When possibility informs reality: Counterfactual thinking as a cue to causality. *Current Directions in Psychological Science*, *8*, 120-123.
- Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, *12*, 49-100.
- Verhagen, A., & Kemmer, S. (1997). Interaction and causation: Causative constructions in modern standard Dutch. *Journal of Pragmatics*, *27*, 61-82.
- Wolff, P., Song, G., & Driscoll, D. (2002). Models of causation and causal verbs. In *Papers from the 37th Meeting of the Chicago Linguistics Society, Main Session, Vol. I*. Chicago: Chicago Linguistics Society.
- Wolff, P. & Song, G. (2001). Models of causation and the semantics of causal verbs. The 42th Annual Meeting of the Psychonomic Society, Orlando, FL.