DE GRUYTER MOUTON
DOI 10.1515/tl-2012-0015 — Theoretical Linguistics 2012; 38(3-4): 237–248

Phillip Wolff

Representing verbs with force vectors

Phillip Wolff: Emory University, Department of Psychology, 321 Psychology Building, 532 N. Kilgo Circle, Atlanta, GA 30322, U.S.A. E-mail: pwolff@emory.edu

1 Introduction

The notion of force may figure in the meaning of several classes of words, including various types of verbs (Copley & Harley, in press; Jackendoff, 1990; Pinker, 1989; Talmy, 1988; Weisgerber, in press; Wolff & Song, 2003; Wolff, Klettke, Ventura, & Song, 2005), prepositions (Zwarts, 2010) and modals (Copley, 2005; Talmy, 1988). In these analyses, it was observed that words encoding force typically display regular patterns of meaning, suggesting an underlying system of distinctions that determines how such meanings are formed (Wolff, Barbey, & Hausknecht, 2010). The most significant step towards characterizing this system was Talmy's (1988) theory of *force dynamics*, which was soon followed by several proposals modifying and extending the theory (e.g., Jackendoff, 1990; Pinker, 1989). In the case of Wolff and colleagues' *dynamics* model (Wolff & Zettergren, 2002; Wolff & Song, 2003; Wolff, 2007), it was shown that Talmy's system of distinctions could be represented by a vector formalism. In the target article of this volume, Warglien, Gärdenfors, and Westera (in press) propose yet another model of how the notion of force is systematically encoded in word meanings. Unlike earlier theories, their *two-vector model* assumes a different set of distinctions than those originally proposed by Talmy (1988). Indeed, it lays the groundwork for a new system of event representation that may help integrate insights from force dynamics with the rest of the lexical semantics literature.

In this review I examine the *two-vector model* with respect to the dynamics model. Both models are recent and make explicit use of vectors in their representations of verb meanings. This comparison will show that the two-vector model and the dynamics model differ with respect to the phenomena they address. I will also discuss several limitations of the two-vector model. In particular, due to the model's single-domain constraint, it seems the theory is unable to represent the meanings of verbs of causation, among others. A second major problem is that the so-called *result vector* introduced in this model seems to over-extend the notion of a vector. Finally, I will discuss how the two-event model implies event representations that are far sparser than those suggested by work in lexical semantics and psychology.

2 Two vector models of verb meaning

2.1 The two-vector model

To facilitate the comparison, I begin with a brief review of the key features of the two-vector and dynamics models. The two-vector model makes a major distinction between (non-linguistic) cognitive representations and linguistic representations. In this model, event representations are cognitive representations. The model holds that event representations specify, at the very least, one *result* vector, one *force* vector, and a central object that undergoes a change, referred to as the *patient*. These three key parts are exemplified in the situation depicted in Figure 1, which shows a person pulling a suitcase over a line. According to the two-vector model, people's cognitive representations of this event will, at the very least, encode the entity being acted on, the patient (the suitcase); the agent's force acting on the patient; and what happens in the event, the result (the suitcase ending up on the other side of the line).

Figure 1 also shows how event representations in the two-vector model can be specified in terms of vectors in space. In this representation, the patient is affected by the force vector from the agent. As with all vectors, the force vector has an origin, a direction, and a magnitude, with the magnitude specified by the length of the vector. The force vector in this particular event representation specifies the cause of the change of location (or state). The event representation also

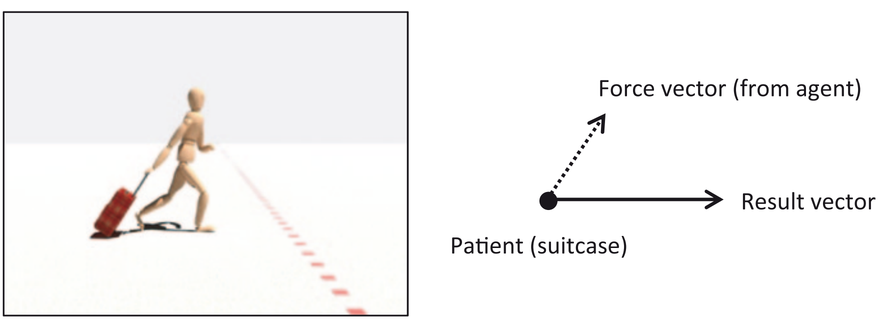

Fig. 1: A situation in which a person pulls a suitcase over across a line, with associated force and result vectors.

(CS4) WDG (155×230mm) DGMetaScience J-2647 TL 38:3-4 pp. 238–248 2647_38_3-4_0015 (p. 238)
PMU:(idp) 3/9/2012 19 September 2012 3:31 PM

includes a result vector, which specifies the patient's movement or displacement through a physical or conceptual space. As noted by Warglien et al., the start and end of the result vector can be viewed as snapshots of the patient at its initial and final positions in the physical or conceptual space.

According to Warglien et al., for many actions, such as lifting or moving a box, the event representation may only encode a single force vector, namely, the one from the agent that acts on the patient. For actions such as walking and swimming, Warglien et al. propose that event representations specify a pattern of interacting forces. According to the two-vector model, event representations can be elaborated on to include additional forces and objects. In particular, event representations may include additional forces acting on the patient, such as the patient's resistance to the change brought about by the agent.

One of the primary goals of the two-vector model is to explain what elements of meaning are denoted by verbs. Verb meanings are hypothesized to be based on only part of an event representation, as stipulated by the *single-domain constraint*. The single-domain constraint holds that a verb can only encode vectors from the same domain. Warglien et al. define a domain as a set of dimensions that are integral rather than separable. Dimensions are integral when having a value on one dimension entails that the other dimension must also have a value (e.g., Nelson, 1993). For example, pitch and loudness are integral because if a sound has a pitch, it must also have loudness. Pitch and hue are not integral because a value on one does not entail that the other dimension also has a value. The single-domain constraint entails that verb meanings will encode only one type of vector in an event representation, either a force or result vector, because force and result vectors come from different domains.

According to the two-vector model, the separation of force and result vectors in the meaning of verbs explains the phenomenon of *manner/result complementarity* (Levin & Rappaport Hovav, 2011; Rappaport Hovav & Levin, 2010), that is, the observation that verbs lexicalize either the manner or result of an event, but not both. Accordingly, manner verbs denote activities often associated with bringing about a change of state e.g., *kick, shake, wipe*) but do not specify any particular change of state, while result verbs denote ending values of a change (e.g., *break, open, fill*) without specifying an associated activity that brings these changes about. Manner/result complementarity is also manifested in expressions of motion, with manner verbs encoding the patterns of movement used to bring about a displacement (e.g., *walk, jog, fly*) and result verbs encoding motion in a specified direction or spatial relation with respect to some point in a trajectory (e.g., *arrive, pass, enter, fall*). Crucially, as with the change of state verbs, motion verbs do not combine manner and result meanings. Again, according to the two-event model, this complementarity occurs because manner and result verbs are

(CS4) WDG (155×230mm) DGMetaScience J-2647 TL 38:3-4 pp. 239–248 2647_38_3-4_0015 (p. 239)
PMU:(idp) 3/9/2012 19 September 2012 3:31 PM

based on different kinds of vectors, force vectors and result vectors, and according to the single-domain constraint, force and result vector will not be lexicalized within a verb.

It should be noted that the single-domain constraint entails that the vast majority of force dynamic meaning representations proposed in Talmy (1988) could not serve as verb meanings. In Talmy's theory, the fundamental dimensions for distinguishing different force dynamic concepts are (a) the intrinsic tendency of the patient for rest or motion, (b) the balance of strengths (i.e., the relative strengths of the agent and patient), and (c) the result of the force interaction, that is, the "overt occurrence" (Talmy, 1988, p. 54) that follows from a force dynamic pattern. Talmy's first two dimensions concern forces while his third concerns result. According to the single-domain constraint, Talmy's first two dimensions could never be lexicalized along with the third dimension because this would entail conflating dimensions from different domains.

2.2 The dynamics model

The dynamics model, which modifies and extends Talmy's (1988) theory of force dynamics, offers an account of how causal concepts are specified in the mind and in word meanings (Wolff & Song, 2003; Wolff, 2007; Wolff, Klettke, Ventura, & Song, 2005; Wolff et al., 2010). The theory holds that causation is specified in terms of configurations of forces that are evaluated with respect to an endstate vector. As with the two-vector model, the forces may be physical, psychological (e.g., intentions), or social (e.g., peer pressure). The theory accounts for several kinds of causal concepts and associated verbs: CAUSE (*cause, force, get, make*), ENABLE (*allows, enable, help, let*), and PREVENT (*block, hinder, keep, prevent*), among others (see Wolff, Klettke, Ventura, & Song, 2005), each associated with a particular configuration of forces. These different concepts are differentiated in terms of three dimensions: 1) the patients' tendency for the endstate, 2) patient-affector concordance, and 3) whether the result is approached. The three dimensions define different possible types of configurations.

A configuration of forces associated with the concept of CAUSE is shown in Figure 2. The configuration is composed of two kinds of vectors: force vectors and position vectors. The position vector specifies the patient's location with respect to the endstate. In a CAUSE configuration, the patient does not have a tendency for the endstate. A patient does not have a tendency for the endstate when it moves itself away from the endstate or simply when it resists movement towards the endstate due to friction. Lack of tendency for the endstate is represented by a patient force, **P**, pointing in a different direction than the position vector, **E**. In a

(CS4) WDG (155×230mm) DGMetaScience J-2647 TL 38:3-4 pp. 240–248 2647_38_3-4_0015 (p. 240)
PMU:(idp) 3/9/2012 19 September 2012 3:31 PM

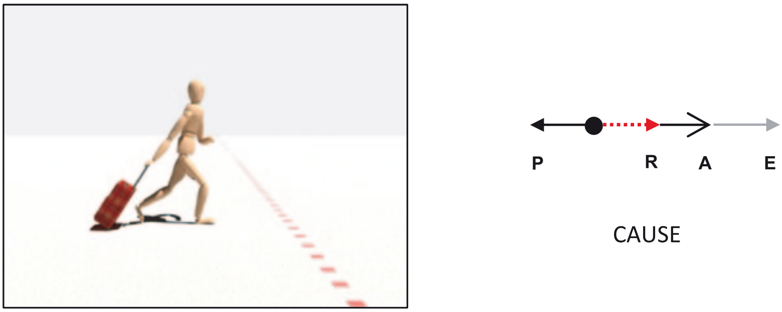

Fig. 2: A free-body diagram (on the right) showing how the dynamics model would specify the force dynamic situation depicted (on the left). **A** = the affector force, **P** = the patient force, **R** = the resultant force, **E** = endstate vector, which is a position vector, not a force.

CAUSE configuration, the affector force, **A**, opposes the patient force, **P**. Lack of concordance is indicated by the affector and patient forces pointing in different directions. Finally, in a CAUSE configuration, the resultant of the **A** and **P** forces, **R**, pushes the patient towards the endstate, which is indicated by the **R** force pointing in the same direction as **E**. Other concepts are given by different configurations of forces. In the case of ENABLE and PREVENT, for example, the patient vector points towards the endstate, capturing the intuition that enabling and preventing require that the patient have a tendency (e.g., intention) for an endstate. CAUSE and ENABLE imply that a result will occur, whereas PREVENT does not. The result distinction is captured, in part, by the resultant vector pointing towards the endstate in the case of CAUSE and ENABLE, but not in the case of PREVENT. The notion of result is captured only *in part* by the dynamics model because configurations of force and an endstate vector can only specify whether a patient is influenced towards a particular endstate, not whether the endstate is actually reached.

3 Comparing the models

3.1 Key differences between the two-vector and dynamics models

The two-vector and dynamics models differ in several ways. First, the models differ in the specification of results. In the two-vector model, a result vector indicates whether a result occurs. In the dynamics model, a resultant vector that points

(CS4) WDG (155×230mm) DGMetaScience J-2647 TL 38:3-4 pp. 241-248 2647_38_3-4_0015 (p. 241)
PMU:(idp) 3/9/2012 19 September 2012 3:31 PM

towards the endstate indicates that the patient will approach a particular state. Barring any changes in the forces, the patient will reach the endstate. Importantly, though, the representations used in the dynamics model do not specify whether the endstate is necessarily reached, just whether it will be approached. A second way the two models differ concerns the kind of phenomena the two models seek to explain. As noted above, the two-vector model offers an account of the difference between manner and result verbs and of manner/result complementarity. The dynamics model, in contrast, focuses on the representation of a particular kind of verb, namely ones associated with the representation of causation and related notions. Moreover, the dynamics model provides a more detailed account of particular verb meanings than is offered by the two-vector model, which, for the most part, does not offer a specific account of any word meanings. Given these differences in emphasis, the models can be viewed as at least partially complementary. However, there is at least one way in which the assumptions and predictions of the two-vector conflict and the dynamics model conflict.

3.2 Challenges for the two-vector model

As discussed earlier, both the two-vector and dynamics models depend heavily on two kinds of vectors. In the two-vector model, the two kinds of vectors are force and result vectors. Due to the single-domain constraint, the two-vector model holds that these two kinds of vectors will not lexicalize in the same verb. In the dynamics model, the two kinds of vectors are force and position vectors. Warglien et al.'s definition of a domain would appear to classify these two kinds of vectors as belonging to different domains, since specifying a force does not entail a relative position and a relative position does not entail a force. Given that force and position vectors belong to different domains, the single-domain constraint entails that representations that include both of these vectors will not lexicalize in a verb. This prediction can be shown to be wrong.

First, causal meanings cannot be represented on the basis of force vectors alone. This fact is perhaps most clearly illustrated with respect to the difference in meaning between CAUSE (e.g., *cause, make, get, force*) and ENABLE (e.g., *allow, enable, help, permit*) verbs. Both types of verbs imply occurrence of a result, so they both would be associated with a result vector. What distinguishes these verbs, according to the dynamics model, is whether the patient has a tendency for the result (Wolff & Song, 2003; Wolff, 2007, Wolff, Barbey, & Hausknecht, 2010). You can *allow* or *enable* someone to do something only if they want to. You *cause* or *force* someone to do something only if they do not want to. As explained in the

(CS4) WDG (155×230mm) DGMetaScience J-2647 TL 38:3-4 pp. 242–248 2647_38_3-4_0015 (p. 242)
PMU:(idp) 3/9/2012 19 September 2012 3:31 PM

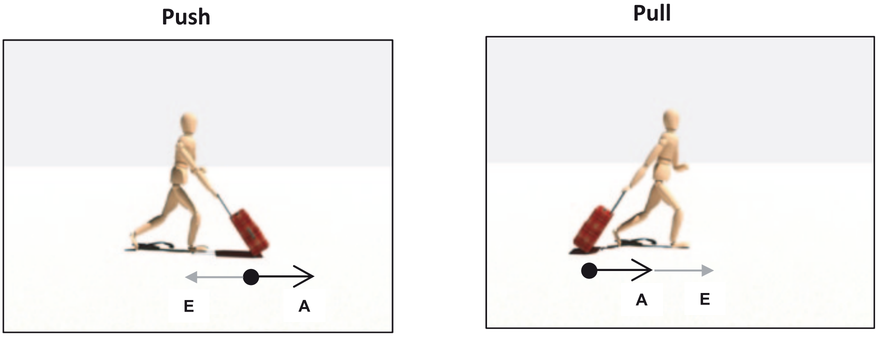

Fig. 3: Scenarios and vectors associated with pushing and pulling. **A** = the affector force, **E** = endstate vector, which starts at the patient and ends at the agent.

dynamics model, specifying a tendency requires knowing the relationship between the force of the patient and some kind of endstate. Such information is well expressed using a position vector. With these distinctions in place, the difference in meaning between CAUSE and ENABLE verbs can be specified. However, this account of the meaning of causal verbs is ruled out by the single-domain constraint of the two-vector model because such meanings conflate vectors from different domains, that is, force and position vectors.

Another informative example is the difference in meaning between the verbs *push* and *pull*. Clearly, the verbs *push* and *pull* denote notions of force, but what makes them opposites? According to Zwarts (2010), the information that differentiates the meanings of these verbs is the direction of the force vectors *relative* to the spatial relation between the agent and the patient. Zwarts' (2010) vector analysis of these verbs can be illustrated with the *push* and *pull* scenarios depicted in Figure 3. In both scenarios, the agent's force on the suitcase is exactly the same respect to the background. What differs is that with pushing, the agent is behind the patient, while with pulling, the agent is in front of the patient. The critical part of this difference can be captured by a position vector that leads from the patient to the agent, in effect, treating the agent's location as an "endstate." In the case of pushing, the affector force, **A**, points in the opposite direction of the endstate vector, whereas in the case of pulling, the affector force points in the same direction as the endstate vector. The patient and resultant forces need not be included in the vector representations because, as is well known, *push* and *pull* verbs do not lexicalize results (Zwarts, 2010). Zwarts' analysis indicates that the meanings of *push* and *pull*, as well as the analyses of several other verb pairs (e.g., *squeeze* and *stretch*, *lean* and *hang*) include in their representations not only forces, but also position vectors that specify spatial directions. In doing

(CS4) WDG (155×230mm) DGMetaScience J-2647 TL 38:3-4 pp. 243–248 2647_38_3-4_0015 (p. 243)
PMU:(idp) 3/9/2012 19 September 2012 3:31 PM

so, these verb meanings also violate the single-domain constraint. In sum, the single-domain constraint appears to be incompatible with several important verb classes.

3.3 The nature of result vectors

In addition to limitations in the representation of verb meanings, the two-vector account is on shaky ground with respect to one of its theoretical constructs. Specifically, result vectors in the two-vector model seem to differ from other sorts of vectors. Vectors, as noted above, specify a magnitude that is associated with their length. In the case of force vectors, a longer vector length implies more force. In the case of position vectors, a longer length implies a greater distance. But in the case of result vectors, what does length imply? More or less result? Unlike with other vector quantities, the notion of magnitude with respect to a result vector is not clear. The problem, it seems, is that achieving a result is a binary distinction, whereas vectors are best suited for representing continuous quantities. The question, then, is whether results are really best modeled as vectors. If they are modeled as vectors, they appear to be vectors of a very different kind than force or position vectors.

3.4 Mental models of events

Unlike most accounts of verb meaning, the two-vector account makes a very clear distinction between representations at the level of (non-linguistic) cognition and those at the level of word meaning. This distinction is important because we know that representations at the level of word meaning are much coarser than those at the level of non-linguistic cognition (for a review see Wolff & Malt, 2010). This difference in granularity is demonstrated in the fact that representations at the level of language are ambiguous in ways that thoughts are not (Pinker, 1994). It is because thoughts are able to make finer representational cuts that people are able to recognize the ambiguities that arise from the coarseness of linguistic representations (e.g., "MINERS REFUSE TO WORK AFTER DEATH"). At the non-linguistic cognitive level, the two-vector model holds that event representations specify, at the very least, one *result* vector, one *force* vector and a central object that undergoes a change, referred to as the *patient*. Warglien et al., acknowledge that people may elaborate on these representations by encoding information about other entities in the scene (e.g., agents, instruments), goals, and patient forces. However, at the same time, they suggest that "[f]or many actions – for

(CS4) WDG (155×230mm) DGMetaScience J-2647 TL 38:3-4 pp. 244–248 2647_38_3-4_0015 (p. 244)
PMU:(idp) 3/9/2012 19 September 2012 3:31 PM

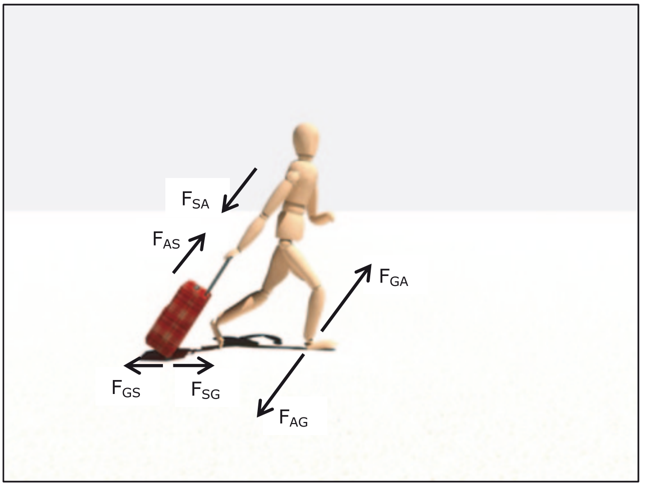

Fig. 4: When an agent (A) walks forward and pulls a suitcase, the following forces are present: $\mathbf{F}_{AG}$ = The agent pushes against the ground, $\mathbf{F}_{GA}$ = The ground pushes back on agent, $\mathbf{F}_{AS}$ = The agent pulls on suitcase, $\mathbf{F}_{SA}$ = The suitcase resists and pulls back on the agent, $\mathbf{F}_{SG}$ = The suitcase wheels push against the ground, $\mathbf{F}_{GS}$ = The friction with the ground resists the suitcase's movement. Note that the magnitude of the action-reaction pairs will, of course, be equal, that is, $||\mathbf{F}_{GA}|| = ||\mathbf{F}_{AG}||$, $||\mathbf{F}_{AS}|| = ||\mathbf{F}_{SA}||$, $||\mathbf{F}_{GS}|| = ||\mathbf{F}_{SG}||$.

example moving and lifting – a single force vector may be sufficient" (Warglien et al., in press, p. 4). While linguistic representations will likely be quite sparse, it should be expected that cognitive representations will be considerably richer, such that they would include more than just a single force vector. To get a sense for what cognitive representations could include, we merely need to consider a more fleshed out representation of the forces in, for example, a moving event like the one shown in Figure 4.

The forces shown in Figure 4 are some of the forces that would be present if an agent walked forward while pulling a suitcase. It is important to note that the action-reaction pairs of forces affect different objects; otherwise, nothing would move. So, for example, when the agent pushes against the ground, the earth accelerates (by an infinitesimal amount). At the same time, the earth pushes back on the agent by the same amount, which, because the agent has much less mass, results in a noticeable acceleration in the agent. Due to the force of gravity (not shown in Figure 4), the agent will not leave the ground, but rather will accelerate in a direction proportional to the horizontal component of the $\mathbf{F}_{GA}$ vector. The net force acting on the suitcase will be the horizontal component of $\mathbf{F}_{AS} - \mathbf{F}_{GS}$; the net force acting on the agent will be the horizontal component of the $\mathbf{F}_{GA}$ – the

(CS4) WDG (155×230mm) DGMetaScience J-2647 TL 38:3-4 pp. 245–248 2647_38_3-4_0015 (p. 245)
PMU:(idp) 3/9/2012 19 September 2012 3:31 PM

horizontal component of $\mathbf{F}_{SA}$. Because the agent moves forward, the magnitude of horizontal component of $\mathbf{F}_{GA}$ must be greater than the magnitude of the horizontal component of $\mathbf{F}_{SA}$ ($\mathbf{F}_{GA} > \mathbf{F}_{SA}$). Because the suitcase moves forward, the force acting on the suitcase in the horizontal direction must be greater than the force of friction ($\mathbf{F}_{AS} > \mathbf{F}_{GS}$).

While it is unlikely that people's mental models of events would encode all of the forces shown in Figure 4, it also seems unlikely that they would encode just one force. Indeed, recent work in perception and psychophysics suggests that people encode at least several of the forces shown in Figure 4 (for a review, see White, 2012). Further, work examining how people describe force dynamic interactions suggests that people must encode more than one force, due to the fact that they readily distinguish causing events from helping events, which entails that they must be attending to forces acting on the patient, such as friction (Wolff, 2007). As a general guideline, people's non-linguistic event representations probably include any of the forces that are regularly denoted in language. Based on this guideline, people's mental models of events are most likely richer than what is implied by Warglien et al. (in press).

4 Conclusions

The two-vector model, with its single-domain constraint, offers an account of how words encoding forces may be related to words encoding results. In addressing this question the two-vector model suggests how vector accounts of verb meaning might be integrated with lexical semantics. However, while the two-vector model is able to account for the phenomenon of manner/result complementarity, it comes at too high a cost, because it predicts the impossibility of certain verb meanings that actually exist.

A number of recent analyses have begun to flesh out what a vector based semantics of verbs might look like. According to Levin and Rappaport Hovav (2005), accounts of verb meaning have been motivated by three major frameworks: 1) the localist approach, in which verb meanings reflect various notions of space (e.g., Jackendoff, 1983, 2) the aspectual approach, which specifies verbs in terms of their temporal properties (Vendler, 1957), and 3) the causal approach, which specifies verbs, in part, in terms of roles in a causal chain (Croft, 1991, to appear). A vector based semantics seems to represent an emerging fourth type of framework. Initially, it could be viewed as part of the causal approach to verbs, but due to the work of Zwarts (2010), Copley (2005), and Warglien et al. (in press), it has become clear that the vector approach can be profitably applied beyond the domain of causation. The work of Warglien et al. (in press) makes a signifi-

(CS4) WDG (155×230mm) DGMetaScience J-2647 TL 38:3-4 pp. 246–248 2647_38_3-4_0015 (p. 246)
PMU:(idp) 3/9/2012 19 September 2012 3:31 PM

cant contribution to the development of this new general framework of word meaning.

References

Copley, Bridget (2005). Ordering and reasoning. In Jon Gajewski, Valentine Hacquard, Bernard Nickel, and Seth Yalcin (Eds.), *New Work on Modality*. New work on modality (MIT Working Papers in Linguistics 52). Department of Linguistics & Philosophy, MIT.

Copley, Bridget, & Harley, Heidi (to appear). An ontology for a force-dynamic treatment of events. In *Forces in Grammatical Structures*, Bridget Copley and Fabienne Martin (eds), Oxford: Oxford University Press.

Croft, William (1991). *Syntactic categories and grammatical relations: The cognitive organization of information*. Chicago, IL: University of Chicago Press.

Croft, William (to appear). *Verbs: Aspect and Argument Structure*. Oxford: Oxford University Press. (Here cited from online version of June 2010.)

Jackendoff, Ray (1983). Semantics and Cognition. Cambridge, MA: MIT Press.

Jackendoff, Ray (1990). Semantic structures. Cambridge, MA: The MIT Press.

Langacker, Ronald W. (1991). Concept, Image, and Symbol: The Cognitive Basis of Grammar. Berlin: Mouton de Gruyter.

Levin, Beth, & Malka Rappaport Hovav (2005). *Argument Realization*. Cambridge: Cambridge University Press.

Levin, Beth, & Malka Rappaport Hovav, M. (to appear). Lexicalized Meaning and Manner/Result Complementarity, In Boban Arsenijević, Berit Gehrke, and Rafael Marín (Eds.), *Subatomic Semantics of Event Predicates*, Springer, Dordrecht.

Nelson, Deborah G. (1993). Processing integral dimensions: the whole view. *Journal of Experimental Psychology: Human Perception and Performance, 19*, 1114–1120.

Pinker, Steve (1989). *Learnability and cognition: The acquisition of argument structure*. Cambridge, MA: MIT Press.

Pinker, S. (1994). *The language instinct*. New York: William Morrow and Company, Inc.

Rappaport Hovav, M., & Levin, B. (2010). Reflections on manner/result complementarity. In Edit Doron, Malka Rappaport Hovav, Ivy Sichel (Eds.), *Syntax, lexical semantics, and event structure*. Oxford University Press, Oxford, UK, 21–38.

Vendler, Zeno (1957). Verbs and times. *Philosophical Review, 56*, pp. 97–121.

Warglien, Massimo, Gärdenfors, Peter, Westera, Matthijs (in press). Event structure, conceptual spaces, and the semantics of verbs. *Theoretical Linguistics*.

Weisgerber, Matthias (in press). Where lexical semantics meets physics: Towards an algorithmic lexical framework for modeling path motion. In Mila Vulchanova and Emile van der Zee (Eds.), *Motion Encoding in Language and Space*, Oxford University Press, Oxford.

White, Peter A. (2012). The experience of force: The role of haptic experience of forces in visual perception of object motion and interactions, mental simulation, and motion-related judgments. *Psychological Bulletin, 138*, 589–615.

Wolff, Phillip, Klettke, Bianca, Tatyana, Ventura, & Song, Grace (2005). Expressing causation in English and other languages. In Woo-Kyoung. Ahn, Robert L. Goldstone, Bradley C. Love, Arthur B. Markman, & Phillip Wolff (Eds.), *Categorization inside and outside the*

(CS4) WDG (155×230mm) DGMetaScience J-2647 TL 38:3-4 pp. 247–248 2647_38_3-4_0015 (p. 247)
PMU:(idp) 3/9/2012 19 September 2012 3:31 PM

laboratory: Essays in honor of Douglas L. Medin (pp. 29–48). Washington, DC: American Psychological Association.

Wolff, Phillip, & Song, Grace (2003). Models of causation and the semantics of causal verbs. *Cognitive Psychology, 47*, 276–332.

Wolff, Phillip & Malt, Barbara, C. (2010). The language-thought interface: An introduction. In Barbara C. Malt & Phillip Wolff (Eds.), *Words and the mind: How words encode human experience* (pp. 3–15). Oxford University Press.

Wolff, Phillip (2007). Representing causation. Representing causation. *Journal of Experimental Psychology: General, 136*, 82–111.

Wolff, Phillip, Barbey, Aron K., & Hausknecht, Matthew (2010). For want of a nail: How absence cause events. *Journal of Experimental Psychology: General, 139*, 191–221.

Wolff, Phillip, & Zettergren, Matthew (2002). A vector model of causal meaning. In W. D. Gray & C. D. Schunn (Eds.), *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society* (pp. 944–949). Mahwah, NJ: Erlbaum.

Zwarts, Joost (1997). Vectors as relative positions: A compositional semantics of modified PPs, *Journal of Semantics, 14*, 57–86.

Zwarts, Joost & Winter, Yoad (2000). Vector space semantics: a model theoretic analysis of locative prepositions, *Journal of Logic, Language and Information, 9*, 171–213.

Zwarts, Joost. (2010). Forceful prepositions. In Vyvyan Evans and Paul Chilton (Eds.), *Language, Cognition and Space: The State of the Art and New Directions*. Equinox Publishing, 193–214.