Original Articles

Psych verbs, the linking problem, and the acquisition of language

Joshua K. Hartshorne [a,b,*], Timothy J. O'Donnell [a,b], Yasutada Sudo [c], Miki Uruwashi [b], Miseon Lee [d], Jesse Snedeker [b]

[a] Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, United States
[b] Department of Psychology, Harvard University, United States
[c] Department of Linguistics and Philosophy, Massachusetts Institute of Technology, United States
[d] Department of English Language and Literature, Hanyang University, Republic of Korea

ABSTRACT

In acquiring language, children must learn to appropriately place the different participants of an event (e.g., causal agent, affected entity) into the correct syntactic positions (e.g., subject, object) so that listeners will know who did what to whom. While many of these mappings can be characterized by broad generalizations, both within and across languages (e.g., semantic agents tend to be mapped onto syntactic subjects), not all verbs fit neatly into these generalizations. One particularly striking example is verbs of psychological state: The experiencer of the state can appear as either the subject (*Agnes fears/hates/loves Bartholomew*) or the direct object (*Agnes frightens/angers/delights Bartholomew*). The present studies explore whether this apparent variability in subject/object mapping may actually result from differences in these verbs' underlying meanings. Specifically, we suggest that verbs like *fear* describe a habitual attitude towards some entity whereas verbs like *frighten* describe an externally caused emotional episode. We find that this distinction systematically characterizes verbs in English, Mandarin, and Korean. This pattern is generalized to novel verbs by adults in English, Japanese, and Russian, and even by English-speaking children who are just beginning to acquire psych verbs. This results support a broad role for systematic mappings between semantics and syntax in language acquisition.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

In learning a language, a child must discover how the different participants in an event are expressed in the sentences that she hears. Otherwise, she'll never be able to explain whether the dog licked her brother (dull) or her brother licked the dog (tattleworthy). Languages signal these roles in a variety of ways including word order, case marking, and grammatical markers on the verb (Dryer & Haspelmath, 2013), and thus solving this *linking problem* necessarily requires learning. But theories of language acquisition differ in their claims about how much is learned and how learning proceeds. A central distinction is between theories which begin with lexically-based generalizations and move towards greater abstraction (Goldberg, 1995, 2006; Tomasello, 1992, 2003; for review see Ambridge & Lieven, 2011, 2014) and theories which invoke broad, abstract links between syntactic and semantic representations from the beginning of language development (Gleitman, 1990; Grimshaw, 1981; Pinker, 1984, 1989).

The relative effectiveness of these two learning strategies depends on what exactly it is that children must learn. One possibility is that human languages are characterized by broad mappings between syntactic roles and semantic roles that apply across predicates of many kinds and which are subject to few if any exceptions (the *systematic mappings hypothesis*). If this is the case, then the learner will benefit from representing sentences in terms of these broad semantic and syntactic roles because doing so will allow her to exploit these mappings to make inferences from meaning to form and from form to meaning. In particular, systematic mappings support syntactic bootstrapping, a process by which children use information about syntax to learn word meanings (Fisher, Gertner, Scott, & Yuan, 2010; Gleitman, 1990; Gleitman, Cassidy, Papafragou, Nappa, & Trueswell, 2005). For example, if a child knows that the subject of a transitive sentence is typically the actor and the object is typically the patient, then she can infer from hearing *The boy gorped the dog* that the new verb must describe the action performed by the boy upon the dog (e.g., petted or fed), rather an action performed by the dog on the boy (e.g., nuzzled or begged). Such regularities are also necessary for semantic bootstrapping, a learning procedure where children use their knowledge of meaning to identify how syntactic functions

* Corresponding author at: Massachusetts Institute of Technology, 77 Massachusetts Avenue, 46-4053H, Cambridge, MA 02139, United States.
E-mail address: jkhartshorne@gmail.com (J.K. Hartshorne).

are marked in their language (Grimshaw, 1981; Pinker, 1984, 1989). For example, if a child knows that agents of caused motion must be realized as the subject of a sentence, she can learn that English marks subjects with word order, that Russian uses inflectional morphology, or that Japanese uses particles, simply by paying attention to how causal agents are marked in the sentences she hears. Critically, on theories like these, broad linking rules are present in languages because every generation of children imposes these kinds of categories on events and sentences (see e.g., Gleitman & Newport, 1995).

In contrast, if linking patterns are arbitrary historical conventions that vary across languages, apply to small sets of verbs, and admit many exceptions, then the strategy above will be counterproductive. Instead, it would be wiser for the child to take a conservative, bottom-up approach to generalization, like those proposed in constructivist theories (e.g., Goldberg, 1995, 2006; Tomasello, 2003). For example, on Tomasello's verb island hypothesis (1992), children initially analyze each predicate as an isolated grammatical island with open argument positions that can be filled with nouns. As more of these lexically-anchored constructions are acquired, children begin to notice the overlap in the semantic functions that are assigned to these fillers, as well as the overlap in their syntactic features (e.g., morphological marking or position relative to the verb). These observations lead them to form broader semantic categories (such as agent and theme), broader syntactic categories (such as verb, subject and object), and generalizations about the relationships between syntax and semantics.

Critically, on both theories adults (and older children) may have broad mappings between syntax and semantics, but the pathway by which they arrive at them is different, as is the degree of systematicity that would be expected both within and across languages.

1.1. Are mappings systematic?

Thus a critical question is whether languages are characterized by broad mappings with few, if any, exceptions (the *systematic mappings hypothesis*). To be precise, the controversy is about the degree of systematicity, not its existence. All theorists recognize that *some* of the syntax-semantics mappings are systematic. For example, in English and many other languages, an agent who causes motion or a change of state in another object is expressed as the subject of an active transitive sentence, rather than the object, regardless of the type of action (Baker, 1988; Croft, 2012; Dowty, 1991; Levin & Rappaport Hovav, 2005; Tenny, 1994). Thus the breaker is the subject of *break*, the drier is the subject of *dry*, the thrower is the subject of *throw* and the liquidator is the subject of *liquidate*. Both adults and toddlers readily extend this pattern to new verbs showing that, for these kinds of events, abstract mapping shapes early acquisition (Dittmar, Abbot-Smith, Lieven, & Tomasello, 2011; Fernandes, Marcus, Di Nubila, & Vouloumanos, 2006; Fisher & Song, 2006; Kline & Demuth, 2014; Marantz, 1982; Noble, Rowland, & Pine, 2011).

The mappings for other kinds events, however, are more controversial, harder to characterize, and arguably more variable both across predicates and across languages (Croft, 2012; Goldberg, 1995; Levin & Rappaport Hovav, 2005). Many verbs with closely related meanings appear to be governed by different linking rules: Agnes's terror of Beatrice might be described with Agnes as the subject (*Agnes feared Beatrice*) or as the direct object (*Beatrice frightened Agnes*). An object moving from Agnes's possession to Beatrice's might be described as *Agnes threw the ball to Beatrice* or as *Agnes threw Beatrice the ball*. A single moment in a game of tag could be described as *Agnes chased Beatrice* or *Beatrice fled from Agnes*.

Taken at face value, such examples call into question the prevalence of broad, systematic mappings from semantics to syntax and the utility of learning procedures that seek to exploit such systematicity (cf. Bowerman, 1988; Boyd & Goldberg, 2011; Braine & Brooks, 1995; Goldberg, 2013). In fact, constructivist theorists have argued that, to acquire this complex mosaic of overlapping and conflicting linking patterns, a learner must begin at the bottom, learning the linking patterns item-by-item and gradually extending them on the basis of the input (Goldberg, 1995, 2006, 2013; Tomasello, 2003). Three lines of evidence lend credence to this account. First, lexically-anchored learning is clearly possible: languages contain idioms and other exceptional mappings (Jackendoff, 2002) and adults readily learn lexically-determined mappings in artificial language studies (Wonnacott, Newport, & Tanenhaus, 2008). Second, in some natural language studies, young children are less likely than older children to generalize novel verbs from one construction to another (see e.g., Akhtar & Tomasello, 1997; Brooks & Tomasello, 1999; Dittmar et al., 2011; Tomasello, 2000), a pattern that is consistent with the claim that syntax-semantics mappings become more abstract over time (but see Fisher, 2002; Kline & Demuth, 2014). Finally, both adults and older children can quickly learn an arbitrary syntax-semantics mapping (one that is unattested in any known language) and generalize it to new verbs (Casenhiser & Goldberg, 2005; Goldberg, Casenhiser, & Sethuraman, 2004), while under these same circumstances five-year olds will only acquire lexically-anchored mappings (Boyd & Goldberg, 2011).

In this paper, we explore a different explanation for these apparent cases of ambiguity and conflict in the linking rules (e.g., chase/flee, fear/frighten). Perhaps the principles that link semantic arguments in an event to syntactic positions in a sentence *are* broad and fully consistent, but we as scientists sometimes fail to see these patterns because we have not correctly characterized the semantic structure. Above we tacitly assumed that if two sentences could refer to the same event, then they had the same semantics, but this assumption is clearly false. Just as the words *dog, canine, pet* and *carnivore* may all be used to refer to the same animal while nonetheless maintaining distinct meanings, sentences may describe the same event while nonetheless picking out different construals – or conceptualizations – of the event (Gentner & Boroditsky, 2001; Gleitman, 1990; Gleitman et al., 2005).

For example, above we noted that in English there are two syntactic structures that are used to describe transfer events:

---

(1) a. Agnes threw/kicked/mailed/shipped the ball to Darpny. (prepositional dative)
     b. Agnes threw/kicked/mailed/shipped Darpny the ball. (double-object dative)

While these two forms are often used to describe similar events, they are argued to express different meanings (Beavers, 2011; Gropen, Pinker, Hollander, Goldberg, & Wilson, 1989; Harley, 2002; Mazurkewich & White, 1984; Hovav & Levin, 2008; Oehrle, 1976). For example, in (1b) Darpny must be a person or organization that is capable of possession, while in (1a) Darpny could simply be a location to which the ball has been sent. This observation has led many theorists to conclude that the two dative constructions express two different semantic structures, or conceptualizations, of transfer events, which are mapped onto two different syntactic forms (Beavers, 2011; Harley, 2002; Hovav & Levin, 2008; Pinker, 1989). Specifically, the *to* dative in (1a) specifies change of location, whereas the double-object data (1b) specifies change of possession. Critically, if the two dative structures have different underlying