Reply to comment

Simulation semantics can revitalize the formalization of meaning: Reply to comments on "Embodied language, best-fit analysis, and formal compositionality"

Jerome Feldman

International Computer Science Institute, 1947 Center St., Berkeley, CA, United States

Yet every sentence is not a proposition; only such are propositions that have in them truth or falsity. Thus a prayer is a sentence, but it is neither true nor false. Let us therefore dismiss all other types of sentences but the proposition, for this last concerns our present inquiry, whereas the investigation of others belongs rather to the study of rhetoric or poetry. Aristotle (De Interpretatione 17a1-8).

I would like to thank the editors of Physics of Life Reviews for soliciting this article and for inviting commentary from a wide range of distinguished linguists and computer scientists. Too often, each subdiscipline is confined to its own silo. Also invited to comment were: Adele Goldberg, Ray Jackendoff, Hans Kamp, Ron Kaplan, Laura Michaelis, and Ivan Sag. As will become obvious, it is very challenging to make progress across paradigm boundaries.

Perhaps the best place to start is with the end of Partee's thoughtful commentary [1]: "Many of Feldman's further goals seem to me to blur distinctions among semantics, pragmatics, world knowledge." Indeed, the target article [2], along with much of linguistics and cognitive science, suggests that the rigid traditional distinctions are counter-productive. The sharp claim here is that a complete treatment of compositionality, as well as practical and psychologically plausible models of language, requires unification. These conflicting perspectives (modular vs. unified) have ignored each other for decades. What is new is that the unified view now seems to be formalized enough that dialog [3] and, hopefully, synthesis are meaningful. In fact, there is a sharp distinction in ECG between *analysis* and *simulation*, depicted in Fig. 1 of the target article [2] and discussed further below.

As has become common, Partee and Steedman [4] equate "formal semantics" and sometimes "semantics" with the methodology that presupposes truth-conditional meaning, traditional modularity, and strong compositionality. As mentioned in the target article, this approach has had remarkable academic success and remains vibrant. There continues to be elegant and insightful work in this tradition, including papers noted by Partee and Steedman. But this

paradigm addresses only part of semantics and is seen as irrelevant by many language researchers and practitioners. Even worse, people in fields like cognitive, functional, and social linguistics often reject all formalization when they find that their concerns cannot be addressed within truth-conditional semantics. It now seems feasible to augment the truth-conditional approach with simulation semantics and a formalization of actions and events like that described in the target article and beyond [5–8].

Steedman [4] is manifestly ambivalent about this; his background is logic and he wants to defend it, but recognizes the limitations. "Logicists have quite a lot to say about inference and entailment, and if they don't have much to say about common sense reasoning on the basis of their semantics, that may be because they have found out the hard way just how difficult it is to make it work." He appears to be not quite ready to consider alternative formalizations of semantics.

Partee, and to some extent Allen [9], seem to assume that the term "compositionality" belongs exclusively to one specialized theory: "The principle has several crucial theory-dependent terms". But the target article [2] is explicitly about the more general, intuitive, notion of the compositionality of language. Most of their comments concern this narrow theory-dependent usage and technical moves that create an internal consistency. It is always possible to define truth-theoretic meaning as functions over possible worlds. Such efforts are discussed in Section 6 of the target article, where I suggest that these are theory-internal formal manipulations that don't address the broader problem of meaning composition.

For example, following the general functional approach, Allen [9] proposes that an adequate treatment of the verb in the Mandarin equivalent of: "You give auntie" is a predicate give2(you, auntie) with no hint on whether the unfilled argument position is a recipient or a theme. His language understanding systems do not, and could not, work with such an impoverished semantics. Even "red hair" has situationally different meanings in a Titian or Modigliani painting, beauty salons, punk rock bands, etc. As a web search will show, there are many contexts with quite different meanings for phrases like "red ball", "red line", etc. [10].

Partee, in another apparent *dictum ex cathedra*, says: "the principle of compositionality in semantics concerns the relation between form and meaning; it does not apply to forms or to meanings separately". Unified (anti-modular?) form-meaning pairs are, of course, the crux of Construction Grammar, which is the carrier for ECG (Section 5). Substituting the word "productivity" for "compositionality" at the conceptual level would have no substantive effect. The fact is that you must have a complex thought (meaning) as a basis for your complex expression (Section 4) and similarly you must be able to grasp a complex concept before you can understand a complex phrase describing it. As Steedman says "after reading the paper, I'm inclined to believe that compositionality is a *necessary* property for anything we would call a concept". ECG and the supporting developments explain (and implement) the form-meaning relation in detail (Section 5).

Steels [11] is generally sympathetic to the approach, but expresses some serious concerns. His first item (shared with Steedman [4]) is the possibility that the evidence for the embodiment of language is just epiphenomenal. It is true that the imaging results alone do not show causal connections, but there is a great deal of behavioral data that does demonstrate the profound interactions between language and action, perception, emotion, etc. This is discussed in Section 2 and the references therein.

His other two problems are the availability of shared context and the link between our X-schema *descriptions* of action and the much more detailed control involved in real animal or robotic motion. Both of these involve the crucial distinction between the SemSpec, which is shared across a language community and simulation (or enactment) that is inherently individual (Fig. 1). Unfortunately, none of the commentators picked up on the central ideas of simulation semantics involving *active representations* and the fine-grained modeling of actions and events, which are the most important and original contributions. I must not have made this sufficiently clear.

The system and theory described in the article focus on communication within a language community, which is assumed to share enough culture as well as grammar. In Steels' work on language evolution [12] or in cross-cultural encounters the problem of finding shared conceptualizations is indeed much more serious. Similarly, our X-schema description of actions (like most of Cognitive Linguistics) is explicitly at the shared conceptual level. The relation between the shared, parameterized X-schema level and the detailed realization of action has been a central concern from the outset [13]. We did some work on automatic mapping from X-schemas to the simulated robot, Jack, but Badler's group at Penn did much more of this with a similar formalism [14]. The ECG X-schemas are explicitly construed as a *parameterization* of the underlying physical system and are postulated as the level used in language; actual simulations use a formalism called CPRM that combines simulation with probabilistic relational models [8].

Allen and Steels make an excellent point that the implemented best-fit analyzer does not currently have an adequate treatment of context. This is discussed in Section 6 of the article, but they both correctly raise additional issues at the discourse level; this is a topic of considerable current interest.

Allen and Steedman also correctly point out that the current implementation shares much with other computational efforts on unification grammar, probabilistic parsing, etc. We have made a concerted effort to be as compatible as possible, both to enhance communication and to exploit relevant insights [15]. But Steedman's dismissive "This too seems fairly standard" misses several key ideas that make it all work. His "logical form like any other" is a deep semantic specification that captures the required conceptual compositionality and supports *simulation semantics*, a radical departure from truth-conditional approaches. Because ECG construction grammar has an embodied semantic pole, many traditional semantic and pragmatic constraints are treated directly in the analyzer. Embodied simulation is not elaborated in this paper, but has been described in many others [8] including Narayanan's 1997 thesis [7]. The best-fit analyzer uses a sophisticated three-component (form, meaning, context) Bayesian metric that enables it to handle hard problems including extraposed constituents and omissions as in the Mandarin example. The same system is able to model detailed psycho-linguistic results on garden path sentences and also is the base for a model of Mandarin grammar acquisition [15]. There is more, but this should be enough to suggest that there are processing lessons to be learned from the current implementation.

The target article is just the beginning of a process. Again, Partee [1] says it best: " … the ambitious goal of giving a compositional account of the "full range of meanings" in human language, one which meets the high standards of formalization found in formal semantics but is built on a theory of "embodied cognition", takes into account the neurobiological substrate of language, and makes use of representations of world knowledge and discourse and situational context as well as linguistic knowledge." The Physics of Life Reviews journal does not currently support continuing commentary, but we have set up the Google group "simulation-semantics" and welcome further discussion.

References

[1] Partee B. Comment on "Embodied language, best fit analysis, and formal compositionality". Phys Life Rev [in this issue].
[2] Feldman J. Embodied language, best-fit analysis, and formal compositionality. Phys Life Rev [in this issue].
[3] Hamm F, Kamp H, van Lambalgen M. There is no opposition between formal and cognitive semantics. Theor Linguist 2006;32:1–40.
[4] Steedman M. Embodied compositionality. Phys Life Rev [in this issue].
[5] Feldman J. From molecule to metaphor: a neural theory of language. Cambridge, MA: MIT Press; 2006.
[6] Feldman J, Narayanan S. Embodied meaning in a neural theory of language. Brain Lang 2004;89:385–92.
[7] Narayanan S. KARMA: Knowledge-based action representations for metaphor and aspect. PhD dissertation. Berkeley, CA: UC Berkeley; 1997.
[8] Narayanan S. Mind changes: A simulation semantics account of counterfactuals [submitted for publication]. http://www.icsi.berkeley.edu/~snarayan/counterfactuals.pdf; 2010.
[9] Allen J. Compositionality, language and intention: Commentary on "Embodied language, best fit analysis, and formal compositionality". Phys Life Rev [in this issue].
[10] Sweetser E. Compositionality and blending. In: Janssen T, Redeker G, editors. Cognitive linguistics: foundations, scope, and methodology. New York: Mouton de Gruyer; 1999. p. 129–62.
[11] Steels L. Does embodiment need to be pushed even further? Comment on "Embodied language, best fit analysis, and formal compositionality". Phys Life Rev [in this issue].
[12] Steels L. Grounding symbols through evolutionary language games. In: Cangelosi A, Parisi D, editors. Simulating the evolution of language. New York: Springer; 2001. p. 221–6.
[13] Bailey D, Feldman J, Narayanan S, Lakoff G. Modeling embodied lexical development. In: Proceedings of the 19th cognitive science society conference; 1997. p. 84–9.
[14] Bindiganavale R, et al. Dynamically altering agent behaviors using natural language instructions. In: Autonomous agents; 2000. p. 293–300.
[15] Bryant J. Constructional analysis. [unpublished PhD dissertation]. Berkeley, CA: UC Berkeley; 2008.