

Is semantics computational?*

MARK STEEDMAN and MATTHEW STONE

1 Introduction

Both formal semantics and cognitive semantics are the source of important insights about language. By developing precise statements of the rules of meaning in fragmentary, abstract languages, formalists have been able to offer perspicuous accounts of how we might come to know such rules and use them to communicate with others. Conversely, by charting the overall landscape of interpretations, cognitivists have documented how closely interpretations draw on the commonsense knowledge that lets us make our way in the world. There is no opposition between these insights. Sooner or later we will have a semantics that responds to both.

However, developing such a semantics is profoundly difficult, because there are certain tensions to be overcome in reconciling the two perspectives. For one thing, the overall landscape of meaning does seem to be characterized by a much richer ontology and more dynamic categories than are exhibited by the fragments typically studied in the formal tradition. One sign of strain is the recent tendency to talk of “procedural”, “non-compositional”, or “computational” semantics, as in Hamm, Kamp and van Lambalgen 2006, hereafter HK&vL. We think such locutions can serve as useful reminders to keep semantics fixed on the central question of how language allows us to share information that some have and others need to get. However, there is some danger that formalists will

* Thanks for support to Stone from Leverhulme Trust, Rutgers University and NSF HLC 0308121, and to Steedman from EU FP6 IST grant PACO-PLUS.

1 merely by put off by an idea that, taken literally, may not be such a good
2 one.

3 In this short article, we want to explore and defend the traditional real-
4 ist view attributed by HK&vL to Lewis among others. In fact, this view
5 offers a well-developed, extremely straightforward and robust account of
6 the relation between semantics and cognition. Moreover, while the realist
7 view has ways of accommodating the representationalist insights of DRT
8 (Lewis 1979; Thomason 1990; Stalnaker 1998), it remains unclear how
9 “computational” semantics can account for the key data for the realist
10 view: cases where we judge interlocutors to be ignorant about aspects of
11 meaning in their native language (Kripke 1972; Putnam 1975; Stalnaker
12 1979; Williamson 1994). This debate about the nature of meaning is deep,
13 substantive, and not likely to be settled soon. However it turns out, formal
14 development as a methodology usually does benefit from separating
15 theories of semantics from theories of processing. We illustrate the point
16 by revisiting the analysis of tense and aspect from HK&vL, and arguing
17 that the non-compositionality they observe is an artifact of the particular
18 representations and reasoning procedures they adopt. Though we phrase
19 our argument as a criticism of a “computational” approach to semantics,
20 we nevertheless hope in many respects to reinforce the moral of HK&vL’s
21 work. Formalist research can and should aim not just for analyses of
22 language structure but explanations of language use. In our view, we are
23 closer than ever to such explanations.

24
25

26 **2 A computational cognitive scientist can be a realist about meaning**

27

28 We regard the representational character of discourse context as a pro-
29 found discovery about linguistic meaning and utterance interpretation.
30 We wholeheartedly endorse this important insight of Discourse Represent-
31 ation Theory (DRT) and the related approaches surveyed by HK&vL.
32 Nevertheless, even if discourse context is representational, it does not
33 follow that discourse context must consist of *mental* representations. We
34 begin by sketching an alternative view. Hold on to your hat.

35 Consider a game of correspondence chess. Two players take it in turns
36 to send each other their moves by email. The moves are expressions in the
37 ordinary notation that chess-players use, spelling out a piece and where to

1 move it: pawn to K4, for example. This is a formal notation, and indeed a
2 computer could parse it, recognize the moves and track the resulting play.
3 Normally, we might expect each player to keep track of the game by
4 moving pieces on a physical chessboard, keeping the board in sync with
5 moves as they are made. But we know this is not necessary. There are
6 people who might forgo the board: experts at chess, or blind players,
7 or players who aim for a contest of memory as well as strategy. These
8 players could use only their imaginations and play a purely virtual game.
9 We suggest that such a game offers a clear example of an abstract, non-
10 mental representation.

11 The state of this chess game is representational in much the same way
12 discourse context is understood to be representational in DRT. Both
13 kinds of representations specify a structured set of variables and their
14 present values. Concretely, in a chess game, there is a variable for whose
15 turn it is. Each piece has a variable for its status; the value stores its cur-
16 rent position or marks the piece as captured. And each square on the
17 board has a variable; its value is the piece that occupies it, or none.
18 (This redundancy is unproblematic. We think of the context as cataloging
19 all the available information, not merely providing a minimally-sufficient
20 record.) The variables that characterize an ongoing chess game, like the
21 discourse referents of a discourse representation structure (DRS), play
22 an indispensable role in describing how events can unfold. Moreover, we
23 face a similar challenge in explaining what these two kinds of representa-
24 tions are. There is no physical array of pieces on a board: we are assum-
25 ing the players have dispensed with such stuff. So, like the discourse refer-
26 ents of a DRS, the variables and values that make up the game state seem
27 to have no autonomous existence in the world.

28 So what kind of representation is the state of the game? Is it some rep-
29 resentation in some player's head? Of course, each player must have a
30 mental representation that tracks the state of the game. However, a little
31 reflection shows that *no* such cognitive structure could *determine* the state
32 of the game. What determines the state of the game is the moves that
33 have been played, and the rules of chess.

34 Let us bring this intuition out. Suppose the players differ in where they
35 think the pieces are. (Playing by memory is challenging.) They may con-
36 test which pieces belong where, but they can resolve the issue by returning
37 to the transcript of moves made and replaying the game. It will turn out

1 that one or the other has made an error. The fact that we can even char-
2 acterize the error as “misrepresentation” shows that the content of the
3 players’ cognitive structures characterize objective circumstances. Players’
4 mental representations can be true or false *of the game*. The game differs
5 from anybody’s representation of it.

6 Indeed, the game can differ from *everybody’s* representation of it. Sup-
7 pose that at a certain stage the two players *both* make the same mistake in
8 tracking the game. At this stage, they mutually suppose, correctly, that
9 each represents the state of the game in that particular way. The two
10 players finish out a brilliant and engrossing contest; there are, not surpris-
11 ingly, some unique turns of strategy. They send the transcript of the game
12 to the newspaper chess column. What happens? Their game will surely be
13 rejected by the editor. It is not chess.

14 It is in this sense that the state of this chess game is an objective, ab-
15 stract social construction. There is a true state of the game. The players
16 do their best to track that state by attending to the moves that are made
17 and working through their knowledge of the rules. To say that chess is an
18 abstract formal system in this sense does not make chess independent of
19 the human mind. The existence of the game surely depends on the fact
20 that there is a community of players that agree on the rules and maintain
21 corresponding mental representations of them. But the view loosens the
22 connection between each specific played-out game and the players’ occur-
23 rent mental representations as they play it. We explain the effect of a par-
24 ticular move in a game not by reference to the mental representations
25 the players update and the computations the players do on the fly, but
26 by reference to the objective state of the game so far and the rules as
27 they prevail in the community.

28 You can read Lewis’s famous paper on *scorekeeping in a language*
29 *game* (1979) as invoking a conception of the state of the conversation
30 that is precisely analogous to the state of a game of chess (actually, base-
31 ball, but you can’t play correspondence baseball). In Lewis’s view, a
32 *conversational score* is an abstract, social representation derived from
33 speakers’ utterances according to the rules of discourse. It is composed
34 of structures that can inform the interpretation of subsequent discourse.
35 The conversational score is thus a natural repository for the collection of
36 discourse referents that structure the context in DRT. In fact, Stalnaker
37 (1998) already makes the point that this kind of discourse context can be

1 regarded as a feature of the real-world social environment in which inter-
2 locutors converse. However, as we have emphasized here, this point is
3 compatible with a fundamentally representationalist view of discourse
4 context. We need not follow Stalnaker (1998) and *reduce* discourse repre-
5 sentations to facts, say, about the form and reference of the linguistic ex-
6 pressions participants have uttered.

7 The distinction is important because it leaves open the possibility of an
8 eclectic and expansive realist understanding of discourse context. For ex-
9 ample, we might discover that discourse context explicitly represents the
10 fact that a name refers to a specific individual, or the fact that a common
11 noun refers to a specific natural kind (Kripke 1972). We might discover
12 that discourse context explicitly represents the standard of strength or
13 precision with which we are to interpret particular vague words (Lewis
14 1979; Williamson 1994). We might discover that discourse context explic-
15 itly represents the agreed purposes of the conversation, the alternative
16 open questions that could be answered next, and the alternative answers
17 that could be given (Thomason 1990; Steedman 2000). The realist view –
18 that the rules of language form an abstract system that, in part, character-
19 ize objective, social representations of discourse context – thus places
20 no constraint on the kinds of information that linguists can use to give
21 rule-governed descriptions of the possible interpretations of utterances in
22 context.

23 The realist view gives not only a coherent account of what language is
24 but enables a coherent, computational account of how we use language.
25 We can know the rules of language: we can maintain mental representa-
26 tions whose content accurately tracks the conventions of our community.
27 We can know the state of the discourse: we can maintain mental represen-
28 tations whose content accurately tracks the real-world social representa-
29 tions of context. And we can draw inferences – computing new mental
30 representations about the interpretation of utterances in context – which
31 faithfully mirror the consequences of the real-world rules in the real-
32 world context. This is just the standard representational theory of mind
33 (Fodor 1987), but it offers an extremely compelling way to narrate the
34 activity of computational mechanisms of language use (Stone 2004).
35 Again, explanations in cognitive science seem no less perspicuous – and
36 no less flexible – when we view mental processes as manipulating struc-
37 tures that *represent* conventions, context, utterances, and their meanings

1 rather than manipulating structures that somehow *constitute* conventions,
2 context, utterances, or their meanings.

3 So far we have just defused arguments against the realist view. But we
4 think there are important reasons to prefer it. They come to the fore as
5 soon as we consider language acquisition. Only a realist can say that lan-
6 guage acquisition is just a case of genuine learning. The story is simple.
7 There are general facts about meaning. The child obtains linguistic experi-
8 ences that give evidence about what these facts are, and thereby arrives
9 inductively at an increasingly precise idea of them. For a computational
10 semantics, by contrast, the best one can hope for is the Quinean project
11 of *radical interpretation*, in which the child attempts holistically to bring
12 the sentences it assents to in line with the sentences used in its community
13 (Quin 1960). In fact, as Quine observed, this empiricist project leads in-
14 evitably to the view that there is in fact no such thing as semantics or
15 knowledge of meaning as distinct from other truths. Realist arguments
16 by Kripke (1972) and Putnam (1975) were indispensable in articulating
17 a principled alternative to this counterintuitive view (Lepore and Stone
18 2006).

19 The central issue is how to describe the linguistic representations of
20 speakers – including, most obviously, language learners – who (a realist
21 would say) are ignorant or incorrect about the reference of words in their
22 native language. Putnam gives himself as an example: he can't tell the
23 difference between elm trees and beech trees. If we think of “computa-
24 tional semantics” as the algorithmic construction of a structured family
25 of sentences that map out what is likely to be true given the information
26 in a sentence, then Putnam's semantics is defective. He has a holistic
27 theory of the world that diverges in systematic ways – regarding elms
28 and beeches – from those of others in his community. Thus, for computa-
29 tional semantics it seems, when it comes to talk of elms and beeches, Put-
30 nam systematically misunderstands the sentences.

31 To a realist, this seems quite a tortured account of a situation that
32 should be described much more straightforwardly. Putnam is a perfectly
33 competent language user, insofar as he intends to use the words *elm* and
34 *beech* with the reference of his community. His linguistic knowledge is
35 partial, so he has a representation that leaves open some details about
36 the reference of those words in English. English, however – that abstract
37 but real social construction – says which kind of tree is which. And

1 Putnam can find out. When he hears and understands utterances like *that's*
2 *an elm*, he gets the evidence he needs to improve his epistemic situation
3 and come to represent the meaning of *elm* in line with his linguistic com-
4 munity (Kripke 1972).

5 The point is not just that this is a clear story. It can be formalized
6 simply and perspicuously (Stalnaker 1979). It can be used to give clear ex-
7 planations that reconcile our intuitions about logic, inference and truth
8 with our intuitions about what we learn when we use sentences of our lan-
9 guage (Williamson 1994). And it's hard to think of any other explanation
10 for how children learn meaning as quickly as they do from such sparse
11 data.

12

13

14 **3 Tense and aspect: A case study**

15

16 HK&vL argue against the realist position. They claim that a nonmono-
17 tonic, noncompositional process of inference is actually necessary to build
18 representations of interpretation that properly account for tense and as-
19 pect. If true, this would undermine the realist view that the grammar of
20 English gives well-defined compositional meanings for tense and aspect.

21 However, it is possible to develop event calculi that offer monotonic
22 reasoning. The semantics of tense and aspect can be spelled out composi-
23 tionally in such calculi. Of course, people may still have to update their
24 representations of events and time nonmonotonically. The world is a non-
25 monotonic place, and under our realist assumptions our mental processes
26 should aim for homomorphic behavior. But this kind of nonmonotonicity
27 remains compatible with our semantics. Once we correctly recognize the
28 grammatically-specified meanings of utterances, we may have to retract
29 some *conclusions* we have drawn from them, but we will never actually
30 have to change the meanings themselves.

31

32

33 **4 Temporal reasoning**

34

35 The “Linear Dynamic” version of the Event Calculus (LDEC Steedman
36 (2002)) is based on dynamic logic (Harel 1984) and linear logic (Girard
37 1987). The combination of these two ingredients gives the system the

1 STRIPS property of modeling change in the world directly via transfor-
 2 mations on representational structures, affording a solution to the compu-
 3 tational form of the frame problem, as well as the representational form
 4 (Shanahan 1997: cf. Fikes and Nilsson 1971), unlike the standard version.

5 Rules in LDEC take the following form, where \rightarrow is linear implica-
 6 tion, α is an action, *Preconditions* are conditions which must hold for the
 7 action to be possible, *Resources* are conditions that cease to hold when α
 8 happens and *Consequences* are conditions that hold after α occurs.

$$9 \quad (1) \quad \{Preconditions\} \wedge Resources \rightarrow [\alpha]Consequences$$

10
 11 The following axiom allows actions to be composed into sequences, or
 12 plans (cf. (Moore, 1980, ch. 3) and Rosenschein (1981)):

$$13 \quad (2) \quad [\alpha][\beta]P \Leftrightarrow [\alpha; \beta]P$$

14
 15 Other control primitives besides seriation are allowed (see 1984: 508). For
 16 example the following LDEC rules represent what a 1–4 month infant
 17 has learned about the breast (simplifying somewhat). First, a breast “af-
 18 fords” sucking, in Gibson’s sense, where \Rightarrow is standard implication:

$$19 \quad (3) \quad breast \Rightarrow affords(suck)$$

20
 21 And the following rule represents the effects of sucking using Kleene +
 22 iteration of a test and an elementary action:

$$23 \quad (4) \quad \{affords(suck)\} \wedge hungry \rightarrow [(hungry?; suck)^+] \neg hungry$$

24
 25 LDEC thus offers a very direct translation of Miller, Galanter and Pri-
 26 bram’s (1960) “TOTE units” or Piaget’s primary (and other) “circular re-
 27 actions” (1936). For example, a slightly older infant may have learned
 28 that wanting to be somewhere affords crawling towards it, and if you
 29 crawl you stop not being there and start being there – simplifying as
 30 usual:

$$31 \quad (5) \quad want(there) \Rightarrow affords(crawl)$$

$$32 \quad (6) \quad \{affords(crawl)\} \wedge \neg there \rightarrow [(\neg there?; crawl)^+] there$$

33
 34 LDEC can be used as the basis for a simple reactive forward chaining
 35 planner. It is all we need to analyze tense and aspect.

36 The linear-dynamic aspects of LDEC embody a solution to both the
 37 representational and computational versions of the ramification frame

1 problem – that change is local. But they do not address the qualification
 2 frame problem – that is, the fact that the best-laid plans go wrong. How
 3 do children learn primary circular reactions in the face of nondeterminis-
 4 tic reality?

5 They might build a vast S4 model of all situations they encounter,
 6 and reason by quantifying over possible worlds. Or they might build a
 7 minimal model based on “inertia worlds” (Dowty 1979), an idea related
 8 to the notion of default, and which Asher (1990) and HK&vL treat
 9 nonmonotonically.

10 More likely, they associate probabilities with rules like (4) and (6),
 11 based on counts of outcomes over those same encountered situations,
 12 and compute directly with probabilities to guide planning. They can then
 13 handle the qualification problem *reactively*, by dealing with failures as
 14 they occur. This can be achieved as long as children have recourse to
 15 rules like those below that are more generally applicable but have such
 16 low probabilities of success that they guide behavior only as a last resort:

17 (7) *affords(bawl)*

18 (8) $\{affords(bawl)\} \wedge \neg happy \multimap [(\neg happy?; bawl)^+] happy$

19
 20 *Mutatis mutandis*, that is roughly our own reaction when we turn the
 21 ignition key of a car and the expected outcome fails to occur.

22
 23
 24 **5 Temporal interpretations**

25
 26 HK&vL interpret the meaning of utterances like those in (9) as instruc-
 27 tions to construct or alter minimal discourse models – canonical repre-
 28 sentations of utterance content – so they explicitly represent *why* the ut-
 29 terances are true.

- 30 (9) a. I have caught the 'flu
 31 b. I am going to Chicago tomorrow
 32 c. John was reaching the other side of the street
 33

34 The *raison d'etre* of such canonical representations is to support com-
 35 putational efficiency by allowing the processor to more easily construct
 36 common patterns of explanation that might otherwise require open-
 37 ended abductive inference. (By the way, we'll stick to the term *canonical*

1 *representation* here because of the unfortunate conflict in terminology be-
2 tween cognitive science and formal semantics, where models are not seen
3 as representations but rather as mathematical abstractions that make the
4 real world itself amenable to precise study.)

5 Constructing these representations is described as being noncomposi-
6 tional and nonmonotonic. Thus, in (9c), a formula concerning a reaching
7 achievement is “coerced” via abduction over world knowledge into a for-
8 mula concerning an activity (say, walking perpendicularly toward the
9 other sidewalk), a “trajectory” including a future goal state (being on
10 the other side of the street), and a realized future state when that goal
11 is attained. If we then hear that John was hit by a truck, then among
12 other additions and changes to the canonical representation, we delete
13 the description of the realized goal state(s) we had previously constructed.
14 Large anterior negative deflections in EEG/ERP measures of brain activ-
15 ity are predicted to accompany such adjustments.

16 Updating knowledge bases in the face of unexpected developments in
17 the world is a long and honorable tradition in artificial intelligence.
18 When the Mars Rover moves its wheels as part of a plan to get to the
19 other side of the canal (which, as is well known, are the nearest thing to
20 roads on Mars), and those wheels slip, so that it falls short of the goal, it
21 updates its estimate of where it is according to a truth-maintenance sys-
22 tem (TMS, Doyle (1979, 1992)), much as I adjust my watch when I hear
23 the time signal for the nine o’clock news: my representation of time
24 says it’s 8:59 but the representation is wrong: the actual socially con-
25 structed time of day is in fact 9:00. What the robot and I are doing is
26 maintaining isomorphism between our representation of the actual world,
27 and the actual world itself, which from our point of view is irritatingly
28 nonmonotonic.

29 However, its not clear that we want to treat representations of linguis-
30 tic content the same way. An utterance is *not* just a description of the
31 actual world, but rather of possible and even counterfactual worlds. A
32 naive semantics for the progressive, according to which the truth of (9c)
33 depends on John’s subsequent state of being on the other side is in danger
34 of falling foul to the original “imperfective paradox” – and predicting
35 that the sentence cannot be true unless John actually did reach the other
36 side of the street. HK&vL’s semantics is not, of course, this naive. The
37 state that is predicated of the reference time by the use of the progressive

1 includes a “trajectory” including the goal state (cf. Shanahan 1997). Even
2 when the activity of moving is marked as *Clipped*, or curtailed by the
3 advent of the truck, there is still enough information to evaluate the truth
4 of statements like “John was reaching the other side when he was hit by
5 the truck,” and “if the truck hadn’t hit him, John would have reached the
6 other side of the street.”

7 But if that is the case, why bother to build a default attained goal state
8 (in the terms of Dowty 1979, the “inertia world”) in the first place? The
9 truth of “John was reaching the other side” can then be directly evaluated
10 with respect to a past situation without reference to any other situation.
11 While it is necessary in due course to represent the fact that the actual
12 course of events either did or did not preempt the default trajectory, this
13 can be done entirely monotonically, by adding descriptions of new event
14 transitions and new states.

15 Indeed, finding canonical representations of specified inertia worlds is
16 frequently impossible. Consider the following examples of the perfect,
17 whose semantics HK&vL also represent in terms of a default past situa-
18 tion when the event in question occurs:

- 19 (10) a. She has caught the ’flu.
20 b. I have forgotten your name.
21 c. We have lost our way.
22 d. I’ve grown accustomed to your face.
23

24 These are all cases in which there either seems to be no single specific past
25 event in question, or where there is no way for the speaker or hearer to
26 locate such an event in time. Indeed the whole point of using the perfect
27 is that when the event happened is immaterial: it’s the consequences that
28 are of interest.

29 By contrast, the true tenses – that is, the past, the present and the futu-
30 rate, as in (11) all involve predication over a reference time which is po-
31 tentially distinct from the time of utterance.

- 32 (11) a. I proclaimed Harry president at ten o’clock last night.
33 b. I proclaim Harry president.
34 c. I proclaim Harry president at ten o’clock tonight.
35

36 We will follow Reichenbach (1947) in calling this temporal referent **R**,
37 distinguishing it from the utterance time **S**, although we will assume **R**

1 represents a length of time rather than a single situation. The past tense
 2 (11a) identifies **R** as preceding **S**. The simple present (11b) identifies the
 3 former as including the latter. The futurate present (11c) identifies a fu-
 4 ture time **R**, of which it predicates properties as factively as the past. It is
 5 as necessary to establish the future temporal referent using phrases like
 6 “at ten o’clock tonight” in (11c) as it is for the past referent in (11a).
 7 This observation implies that the actual realized future is as definite an
 8 element in our realist temporal ontology as the actual realized past, al-
 9 though we inevitably know less about it.

10 It follows that, in the road-crossing scenario, it is simply false to utter
 11 the futurate “#John reaches the other side (soon)” before the truck hit,
 12 just as it is to utter the past “#John reached the other side (then)” after.
 13 If the truck is a surprise, we may have to revise our representations when
 14 we find out it hits. We may make corresponding adjustments about which
 15 utterances we think would be true, and which false. But what we will not
 16 be doing is adjusting our representations of what those utterances would
 17 actually *mean*. That remains as compositional as ever.

18 Nor is it necessary to regard coercion – that is, the transformation in
 19 the context of the progressive construction of an event of one Vendlerian
 20 type, such as the achievement of *reaching the other side of the street*, into
 21 one of another type, such as the activity of *walking perpendicularly to-*
 22 *wards the other side of the street* – as noncompositional.

23 It is easiest to demonstrate this by example, via a compositional gram-
 24 mar fragment for the progressive, pairing categories and interpretations
 25 with the colon operator, presented using the rewrite arrow \rightarrow as a Defi-
 26 nite Clause Grammar (DCG). This notation is simply some syntactic
 27 sugar that turns the definite clauses in a logic program into grammar
 28 rules sensitive to linear order, making the grammar seamlessly compatible
 29 with world knowledge expressed in the LDEC event calculus as con-
 30 straint logic programs.

31

32 (12) $S : vp(np) \rightarrow NP : np VP : vp$
 33 $VP : \lambda x.(coerce(p(x), activity, a) \wedge consequent(p(x), c)$
 34 $\wedge prediction(c) \wedge in_progress(a))@R \wedge tns(R, S)$
 35 $\rightarrow BE + TNS : tns VP_{ING} : p$
 36 $VP_{ING} : \lambda x.p(x, y)$
 37 $\rightarrow V_{ING} : p NP : y$

1 The logical form that this grammar assigns to *John is reaching the other*
 2 *side of the road* (abbreviated *reach(john, other_side)*) is the following:

$$\begin{aligned}
 &3 \text{ (13) } (\textit{coerce}(\textit{reach}(\textit{john}, \textit{other_side}), \textit{activity}, a) \\
 &4 \quad \wedge \textit{consequent}(\textit{reach}(\textit{john}, \textit{other_side}), c) \\
 &5 \quad \wedge \textit{prediction}(c) \wedge \textit{in_progress}(a) @ \mathbf{R} \wedge \mathbf{R} = \mathbf{S} \\
 &6
 \end{aligned}$$

7 This logical form is instantiated by non-grammatical world knowledge
 8 as follows. First, *coerce* coerces anything that is not an activity to be an
 9 activity by finding an accomplishment for which it is the achievement:

$$\begin{aligned}
 &10 \text{ (14) a. } \textit{activity}(p) \Rightarrow \textit{coerce}(p, \textit{activity}, p) \\
 &11 \quad \text{b. } \textit{activity}(p) \wedge \textit{accomplishment}(p, q) \Rightarrow \textit{coerce}(q, \textit{activity}, p) \\
 &12
 \end{aligned}$$

13 Second, specific world knowledge captures the fact that one characteristic
 14 activity that results in reaching a location (abbreviated *reach(x, l)*) is iter-
 15 ated walking towards that location, where *walk(x, l)* is an abbreviation
 16 for $(\neg \textit{at}(x, l)?; \textit{face}(x, l); \textit{step}(x))^+$, that the consequent state of walk-
 17 ing is being *tired*, and that that of *reaching* a location is being *at* that
 18 location.

$$\begin{aligned}
 &19 \text{ (15) a. } \textit{activity}(\textit{walk}(x, l)) \\
 &20 \quad \text{b. } \textit{achievement}(\textit{reach}(x, l)) \\
 &21 \quad \text{c. } \textit{accomplishment}(\textit{walk}(x, l), \textit{reach}(x, l)) \\
 &22 \quad \text{d. } \textit{consequent}(\textit{reach}(x, l), \textit{at}(x, l)) \\
 &23 \quad \text{e. } \textit{consequent}(\textit{walk}(x, l), \textit{tired}(x)) \\
 &24
 \end{aligned}$$

25 On the basis of this general knowledge, the compositionally derived logi-
 26 cal form (13) is instantiated as follows, and will be true just in case John
 27 is walking in that direction with that consequence:

$$\begin{aligned}
 &28 \text{ (16) } (\textit{coerce}(\textit{reach}(\textit{john}, \textit{other_side}), \textit{activity}, \textit{walk}(\textit{john}, \textit{other_side})) \\
 &29 \quad \wedge \textit{consequent}(\textit{reach}(\textit{john}, \textit{other_side}), \textit{at}(\textit{john}, \textit{other_side})) \\
 &30 \quad \wedge \textit{prediction}(\textit{at}(\textit{john}, \textit{other_side}))) \\
 &31 \quad \wedge \textit{in_progress}(\textit{walk}(\textit{john}, \textit{other_side})) @ \mathbf{R} \wedge \mathbf{R} = \mathbf{S} \\
 &32
 \end{aligned}$$

33 The imperfective paradox is avoided: the truth of the proposition is inde-
 34 pendent of whether or not the prediction was fulfilled.

35 The connection between the above semantics for the progressive and
 36 LDEC planning and knowledge about actions is as follows. In order for
 37 (16) to hold, the agent *john* must plan from the goal *at(john, other_side)*

1 using the following knowledge: if you aren't at a place and you aren't
2 walking towards it you can start doing so:

$$3 \quad (17) \quad \neg at(x, l) \wedge \neg in_progress(walk(x, l)) \Rightarrow affords(start(walk(x, l)))$$

5 If you start walking somewhere you are walking somewhere:

$$6 \quad (18) \quad \{affords(start(walk(x, l)))\}$$

$$7 \quad \quad \quad \rightarrow [start(walk(x, l))]in_progress(walk(x, l))$$

9 Walking somewhere and being there affords reaching that place:

$$10 \quad (19) \quad in_progress(walk(x, l)) \wedge at(x, l) \Rightarrow affords(reach(x, l))$$

12 Reaching somewhere means you stop walking and are there:

$$13 \quad (20) \quad \{affords(reach(x, l))\}$$

$$14 \quad \quad \quad \wedge in_progress(walk(x, l)) \rightarrow [reach(x, l)]at(x, l)$$

16 The semantics of the VP rule in the grammar (12) and the definition
17 (14) of coercion imply the following logical form for *John is walking* (cf.
18 (13)), in which $walk(x)$ abbreviates $(step(x))^+$:

$$19 \quad (21) \quad (coerce(walk(x), activity, a) \wedge consequent(walk(x), c)$$

$$20 \quad \quad \quad \wedge in_progress(a) \wedge prediction(c))@R \wedge R = S$$

22 Context and world knowledge now support an instantiation with a as
23 $walk(john)$, with c as $tired(john)$ and with R as S . This process yields a
24 contribution that's true just in case John walking is in progress and the
25 predicted consequence is tiredness:

$$26 \quad (22) \quad (coerce(walk(x), activity, walk(x))$$

$$27 \quad \quad \quad \wedge consequent(walk(x), tired(john))$$

$$28 \quad \quad \quad \wedge in_progress(walk(x)) \wedge prediction(tired(john)))@R \wedge R = S$$

30 is true just in case John is walking with the program of walking and the
31 prediction that the consequences of walking hold.

32 Events involving inanimate objects also have predicted consequences,
33 as in the following logical form for *the door is closing*:

$$34 \quad (23) \quad (coerce(close(door), activity, a) \wedge consequent(close(door), c)$$

$$35 \quad \quad \quad \wedge prediction(c) \wedge in_progress(a))@R \wedge R = S$$

1 Assuming the obvious knowledge about doors, this is instantiated as
2 follows:

3 (24) $(coerce(close(door), activity, swing(door))$
4 $\wedge consequent(close(door), shut(door))$
5 $\wedge prediction(shut(door)) \wedge in_progress(swing(door)))@R \wedge R = S$
6

7 Other categories discussed by HK&vL, such as the perfect, can as we
8 have noted be treated similarly compositionally.

9

10

11 **6 Conclusion**

12

13 Like HK&vL, we envision a science of language that embraces formal
14 methods and cognitive constraints – we strive for this goal in our own
15 research. We too think computation gives important insights into what
16 language is and how language works. And we too think a good scientific
17 theory of language will show that our capacity to learn, understand and
18 produce language relies heavily and closely on cognitive abilities shared
19 with our primate ancestors – in particular, those abilities relating to
20 object-oriented planned action.

21 However, to explain why language can help us get things done at all,
22 it's enough that sentences can be true and usable in principle. Any seman-
23 tic theory offers such an explanation. To explain how language got to be
24 the way it is, we need an argument that really connects language to our
25 characteristic activities – as social animals perceiving, planning, and act-
26 ing in a physical world – and to the ancestral cognitive capacities that
27 underpin those activities. This explanation will appeal to computational
28 principles, but it will appeal equally to empirical facts about our ancestral
29 environment and ancestral biology. We are all a long way from such an
30 account, HK&vL included. We differ methodologically from HK&vL in
31 that we see the key difficulty for this account not in linking semantics to
32 particular theories of reasoning, but in connecting semantics to the real-
33 world settings, social relationships and cognitive architectures which give
34 it its place in nature.

35

36

37

University of Edinburgh
steadman@inf.ed.ac.uk, matthew.stone@rutgers.edu

1 **References**

- 2
- 3 Asher, Nicolas. (1990). "A Default Truth Conditional Semantics for the Progressive." *Linguistics and Philosophy*, 15, 469–598.
- 4 Dowty, David. (1979). *Word Meaning in Montague Grammar*. Dordrecht: Reidel.
- 5 Doyle, Jon. (1979). "A Truth Maintenance System." *Artificial Intelligence*, 12, 231–272.
- 6 —(1992). "Reason Maintenance and Belief Revision: Foundations versus Coherence Theories." In Peter Gardenfors, ed., *Belief Revision*, 29–51. Cambridge: Cambridge University Press.
- 8 Fikes, Richard, and Nils Nilsson. (1971). "STRIPS: a New Approach to the Application of Theorem Proving to Problem Solving." *Artificial Intelligence*, 2, 189–208.
- 9 Fodor, Jerry. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.
- 10 Gibson, James. (1966). *The Senses Considered as Perceptual Systems*. Boston, MA: Houghton-Mifflin Co.
- 11
- 12 Girard, Jean-Yves. (1987). "Linear Logic." *Theoretical Computer Science*, 50, 1–102.
- 13 Hamm, Fritz, Hans Kamp, and Michiel van Lambalgen. (2006). "There is No Opposition between Formal and Cognitive Semantics." This issue.
- 14 Harel, David. (1984). "Dynamic Logic." In Dov Gabbay and F. Guenther, eds., *Handbook of Philosophical Logic*, vol. II, 497–604. Dordrecht: Reidel.
- 15
- 16 Kripke, Saul. (1972). "Naming and Necessity." In Donald Davidson and Gilbert Harman, eds., *Semantics of Natural Language*, 253–355. Dordrecht: Reidel.
- 17
- 18 Lepore, Ernest, and Matthew Stone. (2006). "Logic and Semantic Analysis." In Dale Jacquette, ed., *Handbook of the Philosophy of Logic*. Amsterdam: Elsevier.
- 19
- 20 Lewis, David. (1979). "Scorekeeping in a Language Game." *Journal of Philosophical Logic*, 8, 339–359.
- 21
- 22 Miller, George, Eugene Galanter, and Karl Pribram. (1960). *Plans and the Structure of Behavior*. New York, NY: Henry Holt.
- 23
- 24 Moore, Robert. (1980). *Reasoning about Knowledge and Action*. Ph.D. thesis, MIT, Cambridge MA. Published as TN-191, Menlo Park CA, SRI International.
- 25
- 26 Piaget, Jean. (1936). *La naissance de l'intelligence chez l'enfant*. Paris: Delachaux et Niestle. Translated 1953 as *The origin of Intelligence in the Child*, Routledge and Kegan Paul.
- 27
- 28 Putnam, Hilary. (1975). *Mind, Language and Reality*. Cambridge: Cambridge University Press.
- 29
- 30 Quine, Willard van Ormond. (1960). *Word and Object*. Cambridge, MA: MIT Press.
- 31
- 32 Reichenbach, Hans. (1947). *Elements of Symbolic Logic*. Berkeley, CA: University of California Press.
- 33
- 34 Rosenschein, Stanley. (1981). "Plan Synthesis: a Logical Perspective." In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, 331–337. Vancouver.
- 35
- 36 Shanahan, Murray. (1997). *Solving the Frame Problem*. Cambridge, MA: MIT Press.
- 37
- 38 Stalnaker, Robert. (1979). "Assertion." In Peter Cole, ed., *Syntax and Semantics*, vol. 9. New York: Academic Press.
- 39
- 40 —(1998). "On the representation of context." *Journal of Logic, Language and Information*, 7, 319.
- 41
- 42 Steedman, Mark. (2000). "Information Structure and the Syntax-Phonology Interface." *Linguistic Inquiry*, 34, 649–689.
- 43
- 44 —(2002). "Plans, Affordances, and Combinatory Grammar." *Linguistics and Philosophy*, 25, 723–753.

- 1 — (2005). *The Productions of Time*. Ms., University of Edinburgh, <http://www.inf.ed.ac.uk/~steadman/>.
- 2 Stone, Matthew. (2004). "Intention, Interpretation, and the Computational Structure of Lan-
- 3 guage." *Cognitive Science*, 28, 781–809.
- 4 Thomason, Richmond. (1990). "Accomodation, Meaning, and Implicature." In Philip
- 5 Cohen, Jerry Morgan, and Martha Pollack, eds., *Intentions in Communication*, 325–363.
- 6 Cambridge, MA: MIT Press.
- 7 Williamson, Timothy. (1994). *Vagueness*. Routledge.
- 8
- 9
- 10
- 11
- 12
- 13
- 14
- 15
- 16
- 17
- 18
- 19
- 20
- 21
- 22
- 23
- 24
- 25
- 26
- 27
- 28
- 29
- 30
- 31
- 32
- 33
- 34
- 35
- 36
- 37

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37