Is Semantics Computational?*

Mark Steedman and Matthew Stone
University of Edinburgh
steedman@inf.ed.ac.uk matthew.stone@rutgers.edu

Draft 1.3, December 29, 2005

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

In this short article, we want to explore and defend the traditional realist view attributed by HK&vL to Lewis among others. In fact, this view offers a well-developed, extremely straightforward and robust account of the relation between semantics and cognition. Moreover, while the realist view has ways of accommodating the representationalist insights of DRT (Lewis 1979; Thomason 1990; Stalnaker 1998), it remains unclear how "computational" semantics can account for the key data for the realist view: cases where we judge interlocutors to be ignorant about aspects of meaning in their native language (Kripke 1972; Putnam 1975; Stalnaker 1979; Williamson 1994). This debate about the nature of meaning is deep, substantive, and not likely to be settled soon. However it turns out, formal development as a methodology usually does benefit from separating theories of semantics from theories of processing. We illustrate the point by revisiting the analysis of tense and aspect from HK&vL, and arguing that the non-compositionality

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

they observe is an artifact of the particular representations and reasoning procedures they adopt. Though we phrase our argument as a criticism of a "computational" approach to semantics, we nevertheless hope in many respects to reinforce the moral of HK&vL's work. Formalist research can and should aim not just for analyses of language structure but explanations of language use. In our view, we are closer than ever to such explanations.

2 A computational cognitive scientist can be a realist about meaning

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

Consider a game of correspondence chess. Two players take it in turns to send each other their moves by email. The moves are expressions in the ordinary notation that chess-players use, spelling out a piece and where to move it: pawn to K4, for example. This is a formal notation, and indeed a computer could parse it, recognize the moves and track the resulting play. Normally, we might expect each player to keep track of the game by moving pieces on a physical chessboard, keeping the board in sync with moves as they are made. But we know this is not necessary. There are people who might forgo the board: experts at chess, or blind players, or players who aim for a contest of memory as well as strategy. These players could use only their imaginations and play a purely virtual game. We suggest that such a game offers a clear example of an abstract, non-mental representation.

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

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

Let us bring this intuition out. Suppose the players differ in where they think the pieces are. (Playing by memory is challenging.) They may contest which pieces belong where, but they can resolve the issue by returning to the transcript of moves made and replaying the game. It will turn out that one or the other has made an error. The fact that we can even characterize the error as “misrepresentation” shows that the content of the players’ cognitive structures characterize objective circumstances. Players’ mental representations can be true or false *of the game*. The game differs from anybody’s representation of it.

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

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

You can read Lewis’s famous paper on *scorekeeping in a language game* (1979) as invoking a conception of the state of the conversation that is precisely analogous to the state of a game of chess (actually, baseball, but you can’t play correspondence baseball). In Lewis’s view, a *conversational score* is an abstract, social representation derived from speakers’ utterances according to the rules of discourse. It is composed of structures that can inform the interpretation of subsequent discourse. The conversational score is thus a natural repository for the collection of discourse referents that structure the context in DRT. In fact, Stalnaker (1998) already makes the point that this kind of discourse context can be regarded as a feature of the real-world social environment in which interlocutors converse. However, as we have emphasized here, this point is compatible with a fundamentally representationalist view of discourse context. We need not follow Stalnaker (1998) and *reduce* discourse representations to facts, say, about the form and reference of the linguistic expressions participants have uttered.

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

The realist view gives not only a coherent account of what language is but enables a coherent, computational account of how we use language. We can know the rules of language: we can maintain mental representations whose content accurately tracks the conventions of our community. We can know the state of the discourse: we can maintain mental representations whose content accurately tracks the real-world social representations of context. And we can draw inferences—computing new mental representations about the interpretation of utterances in context—which faithfully mirror the consequences of the real-world rules in the real-world context. This is just the standard representational theory of mind (Fodor 1987), but it offers an extremely compelling way to narrate the activity of computational mechanisms of language use (Stone 2004). Again, explanations in cognitive science seem no less perspicuous—and no less flexible—when we view mental processes as manipulating structures that *represent* conventions, context, utterances, and their meanings rather than manipulating structures that somehow *constitute* conventions, context, utterances, or their meanings.

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

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

To a realist, this seems quite a tortured account of a situation that should be described much more straightforwardly. Putnam is a perfectly competent language user, insofar as he intends to use the words *elm* and *beech* with the reference of his community. His linguistic knowledge is partial, so he has a representation that leaves open some details about the reference of those words in English. English, however—that abstract but real social construction—says which kind of tree is which. And Putnam can find out. When he hears and understands utterances like *that's an elm*, he gets the evidence he needs to improve his epistemic situation and come to represent the meaning of *elm* in line with his linguistic community (Kripke 1972).

The point is not just that this is a clear story. It can be formalized simply and perspicuously

(Stalnaker 1979). It can be used to give clear explanations that reconcile our intuitions about logic, inference and truth with our intuitions about what we learn when we use sentences of our language (Williamson 1994). And it's hard to think of any other explanation for how children learn meaning as quickly as they do from such sparse data.

3 Tense and aspect: A case study

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

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

4 Temporal Reasoning

The "Linear Dynamic" version of the Event Calculus (LDEC Steedman (2002)) is based on dynamic logic (Harel 1984) and linear logic (Girard 1987). The combination of these two ingredients gives the system the STRIPS property of modeling change in the world directly via transformations on representational structures, affording a solution to the computational form of the frame problem, as well as the representational form (Shanahan 1997: cf. Fikes and Nilsson 1971), unlike the standard version.

Rules in LDEC take the following form, where $\multimap$ is linear implication, α is an action, *Preconditions* are conditions which must hold for the action to be possible, *Resources* are conditions that cease to hold when α happens and *Consequences* are conditions that hold after α occurs.

(1) $\{Preconditions\} \wedge Resources \multimap [\alpha]Consequences$

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

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

Other control primitives besides seriation are allowed (see Harel 1984:508). For example the following LDEC rules represent what a 1-4 month infant has learned about the breast (simplifying somewhat). First, a breast "affords" sucking, in Gibson's sense, where $\Rightarrow$ is standard implication:

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

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

(4) $\{\mathit{affords}(\mathit{suck})\} \wedge \mathit{hungry} \multimap [(\mathit{hungry}?;\mathit{suck})^{+}]\neg \mathit{hungry}$

LDEC thus offers a very direct translation of Miller, Galanter and Pribram's (1960) "TOTE units" or Piaget's primary (and other) "circular reactions" (1936). For example, a slightly older infant may have learned that wanting to be somewhere affords crawling towards it, and if you crawl you stop not being there and start being there—simplifying as usual:

(5) $\mathit{want}(\mathit{there}) \Rightarrow \mathit{affords}(\mathit{crawl})$

(6) $\{\mathit{affords}(\mathit{crawl})\} \wedge \neg \mathit{there} \multimap [(\neg \mathit{there}?;\mathit{crawl})^{+}]\mathit{there}$

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

The linear-dynamic aspects of LDEC embody a solution to both the representational and computational versions of the ramification frame problem–that change is local. But they do not address the qualification frame problem—that is, the fact that the best-laid plans go wrong. How do children learn primary circular reactions in the face of nondeterministic reality?

They might build a vast S4 model of all situations they encounter, and reason by quantifying over possible worlds. Or they might build a minimal model based on "inertia worlds" (Dowty 1979), an idea related to the notion of default, and which Asher (1990) and HK&vL treat non-monotonically.

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

(7) $\mathit{affords}(\mathit{bawl})$

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

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

5 Temporal Interpretations

HK&vL interpret the meaning of utterances like those in (9) as instructions to construct or alter minimal discourse models—canonical representations of utterance content—so they explicitly represent *why* the utterances are true.

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

The *raison d'etre* of such canonical representations is to support computational efficiency by allowing the processor to more easily construct common patterns of explanation that might otherwise require open-ended abductive inference. (By the way, we'll stick to the term *canonical representation* here because of the unfortunate conflict in terminology between cognitive science and formal semantics, where models are not seen as representations but rather as mathematical abstractions that make the real world itself amenable to precise study.)

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

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

However, its not clear that we want to treat representations of linguistic content the same way. An utterance is *not* just a description of the actual world, but rather of possible and even counterfactual worlds. A naive semantics for the progressive, according to which the truth of (9c) depends on John's subsequent state of being on the other side is in danger of falling foul to the original "imperfective paradox"—and predicting that the sentence cannot be true unless John actually did reach the other side of the street. HK&vL's semantics is not, of course, this naive. The state that is predicated of the reference time by the use of the progressive includes a "trajectory" including the goal state (cf. Shanahan 1997). Even when the activity of moving is marked as *Clipped*, or curtailed by the advent of the truck, there is still enough information to evaluate the truth of statements like "John was reaching the other side when he was hit by the truck," and "if the truck hadn't hit him, John would have reached the other side of the street."

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

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

(10) a. She has caught the 'flu.
b. I have forgotten your name.
c. We have lost our way.
d. I've grown accustomed to your face.

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

By contrast, the true tenses—that is, the past, the present and the futurate, as in (11) all involve predication over a reference time which is potentially distinct from the time of utterance.

(11) a. I proclaimed Harry president at ten o'clock last night.
b. I proclaim Harry president.
c. I proclaim Harry president at ten o'clock tonight.

We will follow Reichenbach (1947) in calling this temporal referent **R**, distinguishing it from the utterance time **S**, although we will assume **R** represents a length of time rather than a single situation. The past tense (11a) identifies **R** as preceding **S**. The simple present (11b) identifies the former as including the latter. The futurate present (11c) identifies a future time **R**, of which it predicates properties as factively as the past. It is as necessary to establish the future temporal referent using phrases like "at ten o'clock tonight" in (11c) as it is for the past referent in (11a). This observation implies that the actual realized future is as definite an element in our realist temporal ontology as the actual realized past, although we inevitably know less about it.

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

Nor is it necessary to regard coercion—that is, the transformation in the context of the progressive construction of an event of one Vendlerian type, such as the achievement of *reaching the other side of the street*, into one of another type, such as the activity of *walking perpendicularly towards the other side of the street*—as noncompositional.

It is easiest to demonstrate this by example, via a compositional grammar fragment for the progressive, pairing categories and interpretations with the colon operator, presented using the rewrite arrow $\longrightarrow$ as a Definite Clause Grammar (DCG). This notation is simply some syntactic sugar that turns the definite clauses in a logic program into grammar rules sensitive to linear order, making the grammar seamlessly compatible with world knowledge expressed in the LDEC event calculus as constraint logic programs.

(12) $S : vp(np) \longrightarrow \quad NP : np \quad VP : vp$

$VP : \lambda x.(coerce(p(x), activity, a) \wedge consequent(p(x), c)$

$\qquad \wedge\, prediction(c) \wedge in\_progress(a))@\mathbf{R} \wedge tns(\mathbf{R}, \mathbf{S})$

$\qquad \longrightarrow \quad BE + TNS : tns \quad VP_{ING} : p$

$VP_{ING} : \lambda x.p(x, y)$

$\qquad \longrightarrow \quad V_{ING} : p \quad NP : y$

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

(13) $(coerce(reach(john, other\_side), activity, a)$
$\wedge\, consequent(reach(john, other\_side), c)$
$\wedge\, prediction(c) \wedge in\_progress(a)\,@\mathsf{R} \wedge \mathsf{R} = \mathsf{S}$

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

(14) a. $activity(p) \Rightarrow coerce(p, activity, p)$
b. $activity(p) \wedge accomplishment(p, q) \Rightarrow coerce(q, activity, p)$

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

(15) a. $activity(walk(x,l))$
b. $achievement(reach(x,l))$
c. $accomplishment(walk(x,l), reach(x,l))$
d. $consequent(reach(x,l), at(x,l))$
e. $consequent(walk(x,l), tired(x))$

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

(16) $(coerce(reach(john, other\_side), activity, walk(john, other\_side))$
$\wedge\, consequent(reach(john, other\_side), at(john, other\_side))$
$\wedge\, prediction(at(john, other\_side))) \wedge in\_progress(walk(john, other\_side))\,@\mathsf{R} \wedge \mathsf{R} = \mathsf{S}$

The imperfective paradox is avoided: the truth of the proposition is independent of whether or not the prediction was fulfilled.

The connection between the above semantics for the progressive and LDEC planning and knowledge about actions is as follows. In order for (16) to hold, the agent *john* must plan from the goal $at(john, other\_side)$ using the following knowledge: if you aren't at a place and you aren't walking towards it you can start doing so:

(17) $\neg at(x,l) \wedge \neg in\_progress(walk(x,l)) \Rightarrow affords(start(walk(x,l)))$

If you start walking somewhere you are walking somewhere:

(18) $\{affords(start(walk(x,l)))\} \multimap [start(walk(x,l))]in\_progress(walk(x,l))$

Walking somewhere and being there affords reaching that place:

(19) $in\_progress(walk(x,l)) \wedge at(x,l) \Rightarrow affords(reach(x,l))$

Reaching somewhere means you stop walking and are there:

(20) $\{affords(reach(x,l))\} \wedge in\_progress(walk(x,l)) \multimap [reach(x,l)]at(x,l)$

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

(21) $(coerce(walk(x), activity, a) \wedge consequent(walk(x), c)$
$\wedge\, in\_progress(a) \wedge prediction(c))\,@\mathbf{R} \wedge \mathbf{R} = \mathbf{S}$

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

(22) $(coerce(walk(x), activity, walk(x)) \wedge consequent(walk(x), tired(john))$
$\wedge\, in\_progress(walk(x)) \wedge prediction(tired(john)))\,@\mathbf{R} \wedge \mathbf{R} = \mathbf{S}$

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

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

(23) $(coerce(close(door), activity, a) \wedge consequent(close(door), c)$
$\wedge\, prediction(c) \wedge in\_progress(a))\,@\mathbf{R} \wedge \mathbf{R} = \mathbf{S}$

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

(24) $(coerce(close(door), activity, swing(door)) \wedge consequent(close(door), shut(door))$
$\wedge\, prediction(shut(door)) \wedge in\_progress(swing(door)))\,@\mathbf{R} \wedge \mathbf{R} = \mathbf{S}$

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

6 Conclusion

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

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

References

Asher, Nicolas. 1990. "A Default Truth Conditional Semantics for the Progressive." *Linguistics and Philosophy*, 15, 469–598.

Dowty, David. 1979. *Word Meaning in Montague Grammar*. Dordrecht: Reidel.

Doyle, Jon. 1979. “A Truth Maintenance System.” *Artificial Intelligence*, 12, 231–272.

Doyle, Jon. 1992. “Reason Maintenance and Belief Revision: Foundations versus Coherence Theories.” In Peter Gardenfors, ed., *Belief Revision*, 29–51. Cambridge: Cambridge University Press.

Fikes, Richard, and Nils Nilsson. 1971. “STRIPS: a New Approach to the Application of Theorem Proving to Problem Solving.” *Artificial Intelligence*, 2, 189–208.

Fodor, Jerry. 1987. *Psychosemantics*. Cambridge, MA: MIT Press.

Gibson, James. 1966. *The Senses Considered as Perceptual Systems*. Boston, MA: Houghton-Mifflin Co.

Girard, Jean-Yves. 1987. “Linear Logic.” *Theoretical Computer Science*, 50, 1–102.

Hamm, Fritz, Hans Kamp, and Michiel van Lambalgen. 2006. “There is No Opposition between Formal and Cognitive Semantics.” *Theoretical Linguistics*, 32. To appear.

Harel, David. 1984. “Dynamic Logic.” In Dov Gabbay and F. Guenthner, eds., *Handbook of Philosophical Logic*, vol. II, 497–604. Dordrecht: Reidel.

Kripke, Saul. 1972. “Naming and Necessity.” In Donald Davidson and Gilbert Harman, eds., *Semantics of Natural Language*, 253355. Dordrecht: Reidel.

Lepore, Ernest, and Matthew Stone. 2006. “Logic and Semantic Analysis.” In Dale Jacquette, ed., *Handbook of the Philosophy of Logic*. Amsterdam: Elsevier.

Lewis, David. 1979. “Scorekeeping in a Language Game.” *Journal of Philosophical Logic*, 8, 339–359.

Miller, George, Eugene Galanter, and Karl Pribram. 1960. *Plans and the Structure of Behavior*. New York, NY: Henry Holt.

Moore, Robert. 1980. *Reasoning about Knowledge and Action*. Ph.D. thesis, MIT, Cambridge MA. Published as TN-191, Menlo Park CA, SRI International.

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.

Putnam, Hilary. 1975. *Mind, Language and Reality*. Cambridge: Cambridge University Press.

Quine, Willard van Ormond. 1960. *Word and Object*. Cambridge MA: MIT Press.

Reichenbach, Hans. 1947. *Elements of Symbolic Logic*. Berkeley, CA: University of California Press.

Rosenschein, Stanley. 1981. “Plan Synthesis: a Logical Perspective.” In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, 331–337. Vancouver.

Shanahan, Murray. 1997. *Solving the Frame Problem*. Cambridge, MA: MIT Press.

Stalnaker, Robert. 1979. "Assertion." In Peter Cole, ed., *Syntax and Semantics*, vol. 9. New York: Academic Press.

Stalnaker, Robert. 1998. "On the representation of context." *Journal of Logic, Language and Information*, 7, 319.

Steedman, Mark. 2000. "Information Structure and the Syntax-Phonology Interface." *Linguistic Inquiry*, 34, 649–689.

Steedman, Mark. 2002. "Plans, Affordances, and Combinatory Grammar." *Linguistics and Philosophy*, 25, 723–753.

Steedman, Mark. 2005. *The Productions of Time*. Ms., University of Edinburgh, `http://www.inf.ed.ac.uk/~steedman/`.

Stone, Matthew. 2004. "Intention, Interpretion, and the Computational Structure of Language." *Cognitive Science*, 28, 781–809.

Thomason, Richmond. 1990. "Accomodation, Meaning, and Implicature." In Philip Cohen, Jerry Morgan, and Martha Pollack, eds., *Intentions in Communication*, 325–363. Cambridge, MA: MIT Press.

Williamson, Timothy. 1994. *Vagueness*. Routledge.