Representing events and discourse
Comments on Hamm, Kamp and van Lambalgen

TIM FERNANDO

In [HKL00] (henceforth HKL), Hamm, Kamp and van Lambalgen declare "there is no opposition between formal and cognitive semantics," notwithstanding the realist/mentalist divide. That divide separates two sides Jackendoff has (in [Jac96], following Chomsky) labeled *E(xternalized)-semantics*, relating language to a reality independent of speakers, and *I(nternalized)-semantics*, revolving around mental representations and thought. Although formal semanticists have (following David Lewis) traditionally leaned towards E-semantics, it is reasonable to apply formal methods also to I-semantics. This point is made clear in HKL via two computational approaches to natural language semantics, *Discourse Representation Theory* (DRT, [KR93]) and the *Event Calculus* (EC) presented in [LH05]. In this short note, I wish to raise certain questions about EC that can be traced to the applicability of formal methods to E-semantics and I-semantics alike. These opposing orientations suggest different notions of time, event and representation.

1 Are times and events "real" or observed?

The linear order on the real numbers is a firmly established model of the continuum, widely used in connection with space and time. EC follows standard practice in tracking the trajectory of a moving object over a continuous path with real numbers as moments of time. Arbitrarily small movements in space implicate arbitrarily small ticks of a clock. This much is, I think, indisputable.

My worry, however, is that the infinite precision of real numbers runs afoul of the bounded precision of linguistic conventions. Our language

Theoretical Linguistics 32–1 (2006), 57–64 0301–4428/06/0032–0057

discretizes, as van der Does and van Lambalgen report, commenting on [LJ93].

> … it seems to be a pervasive feature of our language faculty to turn analogue features into discrete ones; e.g. the continuous scale of length is partitioned as: *tall, medium, short*. ([dDvL00], page 80)

We can talk about the precise moment of change only if that precision is bounded by discrete approximations of time. It would be absurd to assert that Pat reached the summit of K2 on the first of January 2000, noon GMT *and not* (say) a picosecond earlier. (This is not helped by the imprecision of where exactly the summit of K2 is.) But drawing such impossibly fine distinctions is exactly what the density of the real line forces on a bivalent (true-or-false) semantics. Saying that fluents hold at intervals of real numbers does not help unless we throw uncountably many intervals out, keeping a half-open interval $(a, b]$ only if, for some fixed tolerance $\varepsilon > 0$ and reference point r, the absolute values $|r - a|$ and $|r - b|$ are multiples of ε. This leaves us essentially with the sequence of intervals

$$\ldots, (r - 2\varepsilon, r - \varepsilon], (r - \varepsilon, r], (r, r + \varepsilon], (r + \varepsilon, r + 2\varepsilon], \ldots.$$

How do we pin down what ε and r are? Again we run into the problem that the infinite precision of the reals outstrips language. Perhaps we can do without r and ε, and instead fix any discrete order isomorphic to the integers (freeing ourselves, by abstraction, from the tyranny of the reals).

But then what about the arbitrarily small clock ticks we mentioned earlier in connection with continuous motion? I think one must simply acknowledge an opposition here between, on the one hand, an I-semantic notion of an event as a discretely ordered set of observations, and, on the other hand, an E-semantic notion of an event as progressing continuously over "real" time. Apart from real time, there is a notion of *observation time* that is discretely ordered, with a well-defined notion of next moment, reflecting granularity. For the purposes of natural language semantics, it is useful to regard the I-semantic notions of event and time as prior to the corresponding E-semantic notions. In particular, the continuity of E-semantics can be derived as the limit of finer and finer granularities in I-semantics. (More in section 3 below.) Breakdowns in continuity, as in the case above of Pat reaching the summit, indicate that our language does not always support arbitrary refinements of granularity. In any case, arbitrary refinements are not necessary to engender an impression of continuity; the illusion of motion pictures depends on our inability to detect discontinuity. Our intuitions of continuity from space must be tempered by a no less compelling picture of computation proceeding discretely step-by-step. It is this sequential structure on plans and computations that is crucial to analyzing natural language temporality at a fixed granularity. The density of the reals yields infinitely many points and intervals at which asking if Pat has reached the summit of K2 brings trouble. One is better off constructing time from events, as in the Russell-Kamp construction reviewed in [LH05], assuming the events in question are definable at the same granularity.[1] The strategy, at any rate, is clear: introduce only as many observation times as required to individuate events, passing over real times in silence. By taking observation time to be prior to real time, we sidestep the difficulty in pinning down the aforementioned parameters r and ε.

2 What do "representations" represent?

A second use of real numbers made in [LH05] is to code up events as inputs to constraint logic programs embodying a theory of events. The idea is that meaning lies in an event calculus which can be expressed as constraint logic programs over the reals. There is, of course, an element of arbitrariness in choosing exactly what real number to code up an event. Equating the event (type) of *Pat swimming* with the number 2484 is meaningless without information about the coding or the event calculus into which the input 2484 is fed. Locating meaning in an event calculus where any number of events are thrown in (after suitable coding) buries the essential structure of individual events, making it awkward to state what the meaning of an event is. The obvious question is: is it really necessary to encode an event into such a calculus to say what it is? The alternative is to assign transparent structures to events, on which operations and relations can be defined directly, and some denotational semantic

---

[1] World War II and the tick of an atomic clock are not. Language requires multiple granularities (e.g. [dDvL00]), and it is in refining granularities that perhaps the density of the reals has a place.

account (respecting compositionality) can be spun. If discourses have discourse representation structures (DRSs), why not event representation structures (ERSs) for events? I shall describe a proposal for ERS in the next section. But first, a few more words about the importance of unravelling code.

Significant attention is paid in programming language theory to abstract data types beyond natural numbers, and various abstractions over and above ordinary recursion theory. The aim is to isolate the essence of a notion, and rather than Gödel numbering a syntactic presentation of it, express as straightforwardly as possible what that syntax means, assigning it a denotation free from junk (which r and ε are, with respect to observation times).

3 Events as strings

For an alternative treatment of events, let us start with a simple example. We might picture an event of *rain from dawn to dusk* as $\boxed{\text{rain, dawn}}\boxed{\text{rain}}\boxed{\text{rain, dusk}}$ or $\boxed{\text{rain, dawn}}\boxed{\text{rain}}^k\boxed{\text{rain, dusk}}$ for some $k \geq 2$, where a box encloses a subset of some fixed finite set Φ of temporal formulas[2] to indicate that they all are observed to hold simultaneously. As the number of snapshots between dawn and dusk, each integer $k \geq 1$ determines an event token (observation record), which we can collect in the event type given by the set

$$\boxed{\text{rain, dawn}}\boxed{\text{rain}}^+\boxed{\text{rain, dusk}} = \{\boxed{\text{rain, dawn}}\boxed{\text{rain}}^k\boxed{\text{rain, dusk}} \mid k \geq 1\}$$

of strings over the alphabet $\mathsf{Pow}(\Phi)$ of subsets of Φ. A natural way to compose languages (string sets) is illustrated by

$$\boxed{\text{rain, dawn}}\boxed{\text{rain}}^+\boxed{\text{rain, dusk}} = \boxed{\text{rain}}^+ \ \& \ \boxed{\text{dawn}}\Box^+\boxed{\text{dusk}}$$

where the *superposition* $L\&L'$ of languages L and L' forms the componentwise union of strings of equal length

$$L\&L' = \bigcup_{n \geq 1} \{(\alpha_1 \cup \alpha_1') \cdots (\alpha_n \cup \alpha_n') \mid \alpha_1 \cdots \alpha_n \in L \text{ and } \alpha_1' \cdots \alpha_n' \in L'\}$$

---

[2] These are essentially the uncoded fluents in [LH05]. More on this below.

(and where $\Box = \varnothing$ *qua* symbol, not language). We can then compare information content through a *subsumption* relation $\trianglerighteq$ defined as language containment

$$L \trianglerighteq L' \Leftrightarrow L \subseteq L\&L'$$

not unlike the logical consequence relation on DRSs described in §2.2 of HKL (page 6).

For the inertial reasoning in [LH05], it is useful to designate some formulas φ in Φ as *inertial* in that they persist forward and backward in time unless a force is applied on them. With the help of non-inertial formulas $\mathsf{F}\varphi$ saying 'a force is applied on φ' ([Fer04]), we can express this persistence as implications

$$\boxed{\varphi}\Box \Rightarrow \Box\boxed{\varphi} + \boxed{\mathsf{F}\varphi}\Box \tag{1}$$

$$\Box\boxed{\varphi} \Rightarrow \boxed{\varphi}\Box + \boxed{\mathsf{F}\varphi}\Box \tag{2}$$

where $+$ can be read as disjunction. The precise semantics of (1) and (2) is spelled out in [FN05] (drawing on [BK03]), but an example should suffice to convey the idea. To bring a string such as $\Box\boxed{\varphi}\Box$ in compliance with (1) and (2), we might let φ flow forward and backward for the *inertial completion* $\boxed{\varphi}\boxed{\varphi}\boxed{\varphi}$, minimizing forces at the cost of maximizing inertial flow. Or we can *f*reeze the flow with forces in the *explanation* $\boxed{\mathsf{F}\varphi}\boxed{\varphi, \mathsf{F}\varphi}\Box$. It turns out that if a language L is regular (i.e. accepted by a finite automaton) then so are its inertial completions $ic(L)$ and explanations ([Fer04]). Moreover, if L and L' are regular, so are $L\&L'$ and the language $L_\Box \supseteq L$ consisting of strings in L with any number of leading and trailing $\Box$'s added or stripped off. (E.g. $(\Box\boxed{\varphi})_\Box = \Box^*\boxed{\varphi}\Box^*$.) Hence, one may apply one's favourite finite-state toolkit (e.g. [BK03]) to compute inertial entailments

$$L \vdash_{ic} L' \Leftrightarrow ic(L) \trianglerighteq L'_\Box.$$

that extend the explicit entailments in $\trianglerighteq$. The relation $\vdash_{ic}$ is but one instance of entailments $\vdash_{\mathscr{E}}$ induced by a map $\mathscr{E}$ (such as ic) on languages L incorporating background assumptions $\mathscr{E}(L) \trianglerighteq L$. Just as $ic(L)$ is derived from the implications (1) and (2) above (by taking the first disjunct), so too can we derive maps $\mathscr{E}$ from implications such as

$$\boxed{\varphi \wedge \psi} \Rightarrow \boxed{\varphi, \psi}$$
$$\Box\boxed{previous(\varphi)} \Rightarrow \boxed{\varphi}\Box$$
$$\boxed{\text{dawn}}\Box^{+}\boxed{\text{dusk}} \Rightarrow \Box^{+}\boxed{\text{noon}}\Box^{+}.$$

The maps $\mathscr{E}$ allow us to smuggle into $\vdash_{\mathscr{E}}$ a theory of events, the vacuous identity map *id* yielding a relation $\vdash_{id}$ for mereology.

What about the event theory in [LH05]? We can picture $HoldsAt(f, t)$, for 'the fluent f holds at time t,' as the string $\boxed{f, \mathsf{time}(t)}$, and $Happens(e, t)$, for 'the event e happens at t,' as the language

$$\mathscr{L}(e) \;\&\; \Box^{*}\boxed{\mathsf{time}(t)}$$

assuming $f, \mathsf{time}(t) \in \Phi$ and e is reducible to a language $\mathscr{L}(e) \subseteq \mathsf{Pow}(\Phi)^{+}$. This treatment preserves the sortal difference between event types and fluents in [LH05], although it is easy enough to turn the fluent f into the language $L = \boxed{f}^{+}$ that is *stative* in that it persists forward and backward, $L = ic(L_{\Box})$ [Fer04] (bringing the distinction between *HoldsAt* and *Happens* in line with that between Hold and Culminate in [Par90]). I believe the completions in [LH05] minimizing events can be matched by the inertial completions above operating on fluents $\mathsf{F}f$ that clip otherwise persistent inertial fluents f.

As for continuous change, we can track the progress of a transition $\boxed{\neg\varphi}\boxed{\varphi}$ by a language

$$L[\varphi\text{-deg}] = \boxed{\varphi\text{-deg}(0)}\boxed{\varphi\text{-deg}_{\uparrow}}^{+}\boxed{\varphi\text{-deg}(1)}$$

built from a degree formula φ-deg(x) such that φ-deg(1) is equivalent to φ, while φ-deg$_{\uparrow}$ records an increment in φ-deg short of φ

$$(\exists x < 1)\; \varphi\text{-deg}(x) \wedge (\exists y < x)\; previous(\varphi\text{-deg}(y)).$$

For instance, for $\varphi =$ *Pat swim a mile since time t*, we might equate the degree x with the fraction x of a mile Pat has swum since t. For a cinematic impression of continuity, we should be careful not to pick a string from $L[\varphi$-deg$]$ that is too short. But neither should it be too long, lest we place impossible demands on our camera, as in the K2 example above. Rather than insisting on a particular string from $L[\varphi$-deg$]$, it is enough that we carry a *set* of possible choices along.

4 Beyond events: DRSs as types of proofs

Turning very briefly to the step from events to DRSs, I propose that a DRT event term e denotes an event token (string), as opposed to an event type (in [LH05] or elsewhere), and that we read the predication $e : arrive(m)$ in §4.1 of HKL (page 23) as e belongs to a suitable language (event type) $\mathscr{L}(arrive(m)) \subseteq \mathsf{Pow}(\Phi)^{+}$. In general, strings in $\mathsf{Pow}(\Phi)^{+}$ serve the needs of "subatomic semantics" ([Par90]), by which Parsons means "the study of those 'formulas in English' that are treated as atomic formulas in most logical investigations." But more complicated DRS-conditions and DRSs (describing, for instance, infinite stretches of time or universal quantification) require types whose tokens go beyond strings in $\mathsf{Pow}(\Phi)^{+}$.[3] There is a well-known theory ([ML84]) of these types that has been applied to natural language semantics (e.g. [Ran94, Coo05]), treating a formula as the type of its proofs. A DRS with a discourse referent can be treated as an existentially quantified formula, a proof of which provides a witness corresponding to the discourse referent. Events expressed as discourse referents in DRT fit in with this treatment inasmuch as the strings associated with a DRS condition can be construed as proofs of that condition. That is, events as strings provide the tokens of subatomic types within a uniform semantics for DRSs and events. I hope to spell out the details elsewhere.

Acknowledgment

My thanks to Fritz Hamm, Hans Kamp and Michiel van Lambalgen for their very stimulating work.

References

[BK03] K. R. Beesley and L. Karttunen. (2003). *Finite State Morphology*. CSLI Publications, Stanford.

---

[3] The step up in syntactic complexity here mirrors a basic obstacle to the marriage of DRT and constraint logic programming proposed in HKL: "the desirable computational properties of the [latter] depend crucially on specific restrictions imposed on its syntax" (page 22).

[Coo05] R. Cooper. (2005). Records and record types in semantic theory. *Journal of Logic and Computation*, 15(2):99–112.

[dDvL00] J. van der Does and M. van Lambalgen. (2000). A logic of vision. *Linguistics and Philosophy*, 23:1–92.

[Fer04] T. Fernando. (2004). Inertia in temporal modification. In *Proc. Semantics and Linguistic Theory XIV*, pages 56–73. Cornell Linguistics Circle Publications.

[FN05] T. Fernando and R. Nairn. (2005). Entailments in finite-state temporality. In *Proc. 6th International Workshop on Computational Semantics*, pages 128–138. Tilburg University.

[HKL00] F. Hamm, H. Kamp, and M. van Lambalgen. There is no opposition between formal and cognitive semantics. This issue.

[Jac96] R. Jackendoff. (1996). Semantics and cognition. In S. Lappin, editor, *The Handbook of Contemporary Semantic Theory*, pages 539–559. Blackwell, Oxford.

[KR93] H. Kamp and U. Reyle. (1993). *From Discourse to Logic*. Kluwer Academic Publishers, Dordrecht.

[LH05] M. van Lambalgen and F. Hamm. (2005). *The Proper Treatment of Events*. Blackwell.

[LJ93] B. Landau and R. Jackendoff. (1993). “What” and “where” in language and spatial cognition. *Behavioral and Brain Sciences*, 16:217–265.

[ML84] P. Martin-Löf. (1980). *Intuitionistic Type Theory*. Bibliopolis, Napoli, 1984. Notes by Giovanni Sambin of a series of lectures given in Padua, June.

[Par90] T. Parsons. (1990). *Events in the Semantics of English: A Study in Subatomic Semantics*. MIT Press, Cambridge, MA.

[Ran94] A. Ranta. (1994). *Type-Theoretical Grammar*. Oxford University Press, Oxford.