Steedman's Temporality Proposal and Finite Automata

Tim Fernando*

Trinity College Dublin, Ireland

Abstract. The proposal from Steedman 2005 that "the formal devices" required for temporality in linguistic semantics "are those related to representation of causality and goal-directed action" is developed using finite automata, implicit in which are notions of causality (labeled transitions) and goal-directed action (final/accepting states). A bounded granularity is fixed through a finite alphabet, the temporality of which is given special attention and applied to causal models. A notion of a string in compliance with a causal model is defined such that for a finite causal model M, the set of strings compliant with M is a regular language.

Keywords: Temporality, causal model, finite automata.

1 Introduction

In a wide-ranging study, Steedman (2005) argues that "the so-called temporal semantics of natural language is not primarily to do with time at all" (as given say, by the real line $\mathbb{R}$), and proposes rather that "the formal devices we need are those related to representation of causality and goal-directed action." Finite automata are simple candidates for such devices, with notions of causality and goal-directed action implicit in an automaton's transition table and accepting (final) states. Stepping back from automata, the strings which such automata may or may not accept are employed in Fernando 2004, 2008 to represent events of various kinds (for linguistic semantics). But does this step (back) not trivialize the notions of causality and goal-directed action offered by finite automata? The present paper is an attempt to bring these notions to the fore, and more generally, to explain why Steedman's proposal – henceforth ST – is interesting, and how it might be implemented through finite-state methods – ST_{fa}.

Motivation for ST can be found from discourse coherence, (1), down to tense and aspect, (2).

(1) Max fell. Mary pushed him.
(2) John [*has] left, but is back.

In (1), the push is most naturally understood as preceding the fall because (1) suggests a causal connection, while in (2), the difference between the simple past and the present perfect (*has left*) is that under the latter, the result of John's

* I thank my anonymous referees for suggesting changes to a previous draft.

M. Aloni et al. (Eds.): Amsterdam Colloquium 2011, LNCS 7218, pp. 301–310, 2012.

departure persists through (2)'s utterance (incompatible with him back). In (1) and (2), time and its modelling by $\mathbb{R}$ are secondary to the changes or, as the case may be, non-changes that are communicated. Put crudely under ST_{fa}, time is the result of running automata. A more moderate position is that time is conceived in ST to be relative (as opposed to absolute or independent), the raison d'être of which is to place some set E of events and states in some order. Fleshed out according to Russell and Wiener (e.g. Lück 2006), this relative conception of time can be brought in line with the ST_{fa}-view of time as runtime (Fernando 2011a). Apart from their runs, however, do the automata merit a place in semantics? Consider the habitual (3a), in contrast to the episodic (3b).

(3) a. Tess eats dal.
    b. Tess is eating dal.

Not only can we assert (3a) and at the same time deny (3b), it is not entirely clear that the truth of (3a) at a world and time can be reduced to its episodic instances at that world (even if we may agree on some minimal constraints between them). Under the *rules-and-regulations* view defended in Carlson 1995, generic sentences such as (3a) are true at a world and time only if at that world and time, "structures" exist that are not "the episodic instances but rather the causal forces behind those instances" (page 225). Can automata serve as such structures? An important test is the notion of a causal model, as described in Pearl 2009. How ST_{fa} might relate to causal models is explored in the next section,[1] which puts into context the development of ST_{fa} sketched in the subsequent sections.

2 Causal Models

A *causal model* $M = \langle \mathcal{U}, \mathcal{V}, \{F_X\}_{X \in \mathcal{V}} \rangle$ consists of a set $\mathcal{U}$ of *background variables*, a disjoint set $\mathcal{V}$ of *endogenous variables*, and for each endogenous variable $X \in \mathcal{V}$, a function F_X returning the value of X, given values for the other variables (in $\mathcal{U} \cup \mathcal{V} - \{X\}$). For $Y \in \mathcal{U} \cup \mathcal{V} - \{X\}$, we say X *depends on* Y if the value of Y may matter to the value of X in that there are two settings $\boldsymbol{v}$ and $\boldsymbol{v'}$ of values for variables in $\mathcal{U} \cup \mathcal{V} - \{X\}$ differing only on Y such that $F_X(\boldsymbol{v}) \neq F_X(\boldsymbol{v'})$. The *causal network* $\mathcal{G}(M)$ *of* M is the directed graph with the set $\mathcal{U} \cup \mathcal{V}$ of nodes and an edge from Y to X precisely if $X \in \mathcal{V}$ and X depends on Y. The intuition is that F_X describes a causal mechanism for X, and edges in $\mathcal{G}(M)$ represent (immediate) causal dependence. An example where the variables are Boolean (ranging over two values 1 and 0 for truth and falsity, respectively) is the shooting squad scenario from Schulz 2011, with four variables

> C for *the court orders the execution*, O for *the officer gives the signal*, R for *the rifleman shoots*, and P for *the prisoner dies* (page 248)

and $\mathcal{U} = \{C\}$, $\mathcal{V} = \{O, R, P\}$, causal network pictured by

$$\boxed{C} \rightarrow \boxed{O} \rightarrow \boxed{R} \rightarrow \boxed{P}$$

[1] I am indebted to the referee who set this task, with pointers that got me on my way.

and F_O returning the value of C, F_R returning the value of O, and F_P returning the value of R. But suppose

(∗) the court issues its execution orders *after* the prisoner dies

or (for the non-veridical case of "the prisoner dies before the order")

(†) the prisoner dies even though the court never orders his execution.

One option is to sharpen our reading of the variables C, O, R and P by associating times $\tau(C), \tau(O), \tau(R)$ and $\tau(P)$ with them such that (at the very least) whenever the causal network has an edge from Y to X, $\tau(X)$ is understood not to precede $\tau(Y)$. Exactly what times are (points or intervals?) or what precedence is, let us put aside for now, whilst heeding the point made in Halpern and Pearl 2005 (page 849) that

> it is always possible to timestamp events to impose an ordering on variables and thus construct a recursive model corresponding to a story

where a model M is said to be *recursive* if its causal network $\mathcal{G}(M)$ is acyclic. The idea for (∗) is that assuming the prisoner dies at $\tau(P)$ (making $P = 1$), the court's orders (afterwards) fall outside $\tau(C)$ (leaving $C = 0$), and rather than being in spurious compliance with the causal model, (∗) becomes an instance of an *external intervention modifying the causal model* (an eventuality fully accounted for in Pearl 2009). So too would (†), assuming the court orders clemency at $\tau(C)$ and the prisoner dies at $\tau(P)$. But what if the court never issues an order (one way or the other)? How are we to define $\tau(C)$? Were the court to order execution a moment after $\tau(C)$ and the prisoner then to die as a result, C would be worthless. Once the prisoner dies, however, $\tau(C)$ may as well be over. But what if the prisoner does not die? What if nothing happens, and the officer, rifleman, and prisoner are left waiting? Is $\tau(C)$ to go on indefinitely, along with $\tau(O)$, $\tau(R)$ and $\tau(P)$? Common sense suggests $\tau(X)$ is a bounded interval that may vary according to the value X takes, and the values of other variables Y and their times $\tau(Y)$.

It is noteworthy that causal models typically leave time out, exceptions being "dynamic" causal models where "variables are time indexed" (Halpern and Pearl 2005, page 863), with integer subscripts on variables distinguishing values pertaining to different times.[2] It is not obvious what the shooting squad example would gain by attaching subscripts on the variables C, O, R, P. And although

[2] This practice appears to run counter to Hiddleston 2005, footnote 5, which reads

> I take the variables to represent properties in the first instance. But I treat an event as an object's having a property at a time, so the variables can equally represent a particular case (events) or a repeatable set-up (properties). The variables must represent properties in the first instance if the graphs are capable of representing multiple situations: individual events (such as Janes being bitten by the snake at a given time and place) occur only once.

Hiddleston repeats this point in page 648 (and in another journal that same year).

the questions above about $\tau(X)$ are designed to clarify what a variable X represents, their answers come at a price: the loss of simplicity is accompanied by a loss of generality, making the causal model less interesting. Can we keep the shooting squad uncluttered by $\tau(X)$? Its causal network above neatly depicts a sequence of four events $C = 1$, $O = 1$, $R = 1$ and $P = 1$ of ordering, signalling, shooting and dying. But the denial $C = 0$ that the event $C = 1$ happens is, under a common view of aspect, a non-event: a state holding at *all* subintervals of an interval, rather than an event happening at *some* subinterval.

We can bring out the distinction here between events and states, forming a string $\alpha_1\alpha_2 \cdots \alpha_n$ of sets α_i of *fluents* (i.e., temporal propositions) to describe a time at which every fluent in α_1 holds, followed by one at which every fluent in α_2 holds, and so on through α_n. Using fluents $\dot{X}$ for the equation $X = 1$, the string $\boxed{\dot{C}}\boxed{\dot{O}}\boxed{\dot{R}}\boxed{\dot{P}}$ of length 4 records the sequence of events $C = 1$, $O = 1$, $R = 1$ and $P = 1$ at four successive times. And using fluents $\neg X$ for the equation $X = 0$, the string $\boxed{\neg C}\boxed{\neg O}\boxed{\neg R}\boxed{\neg P}$ also of length 4 describes a time at which $C = 0$ followed by a time at which $O = 0$, and then $R = 0$ and finally $P = 0$. Now, the difference between an event such as $C = 1$ and a state such as $C = 0$ is that whereas the fluent $\dot{C}$ expresses a *force* for change, the fluent $\neg C$ for the latter can be assumed to be *inertial*, where

(I) an inertial fluent persists forward and backward unless forced otherwise

(Fernando 2008). The upshot for $\boxed{\neg C}\boxed{\neg O}\boxed{\neg R}\boxed{\neg P}$ is that in the absence of external forces, (I) entails that the inertial fluents $\neg C$, $\neg O$, $\neg R$ and $\neg P$ flow forward and backward in $\boxed{\neg C}\boxed{\neg O}\boxed{\neg R}\boxed{\neg P}$ to yield the string

$$\hat{s} := \boxed{\neg C, \neg O, \neg R, \neg P}\boxed{\neg C, \neg O, \neg R, \neg P}\boxed{\neg C, \neg O, \neg R, \neg P}\boxed{\neg C, \neg O, \neg R, \neg P}.$$

Furthermore, drawing on a relative conception of time as change, the string $\hat{s}$ *block-reduces* to $\boxed{\neg C, \neg O, \neg R, \neg P}$, the one symbol in $\hat{s}$. More on time-as-change and block reduction in the next section. For now, suffice it to say that in case $C = O = R = P = 0$, we are left with a block reduced string of length one, in contrast to the string $\boxed{\dot{C}}\boxed{\dot{O}}\boxed{\dot{R}}\boxed{\dot{P}}$ of length 4 for $C = O = R = P = 1$.[3] Other combinations of values for the variables are possible, but these arise from interventions to the causal model referred to by the "unless"-clause in (I) above. (I) arose in Fernando 2008 to account for the so-called *Imperfective Paradox* (e.g. Dowty 1979), the point being that interruptions to the progressive are, along with interventions to the causal model and *finks*, *masks* or *antidotes* in dispositions (e.g. Bird *ta*), external forces that may interfere with a causal picture (conceived in isolation from other pictures). Rather than take for granted some notion of

[3] A block-reduced string of length 4 can also be associated with $C = O = R = P = 0$ if, for instance, we add four non-inertial fluents marking the times at which the variables are determined separately to have value 0. As should become clear below, the conception time-as-change is relative to the fluents allowed to appear in boxes.

possible world, and interpret an event $X = x$ as a set of possible worlds (i.e., a proposition, as in footnote 6 of Halpern and Pearl 2005, and footnote 5 of Schulz 2011), the approach pursued under ST_{fa} below is to build up a temporal realm of occurrences from causal mechanisms that execute alongside other (possibly interfering) causal mechanisms. In ST_{fa}, a causal mechanism is a finite automaton that operates in discrete steps to induce a discrete notion of time. That notion of time is subject to refinement insofar as other automata may run before, during or after it (or indeed instead of parts of it).

To keep matters simple first, however, let us focus on a single causal model M that makes no mention of time. If we can restrict the causal model's set of variables to a finite set, and the values these variables take to a finite set (as we can with the shooting squad), then the set of strings that comply (atemporally) with M is accepted by a finite automaton (i.e. a regular language). To make all this precise, more notation is necessary. Given a causal model M in which each variable X takes values from a set $\mathsf{Val}(X)$, let us let us collect the possible variable-value pairs under M in

$$\mathsf{vv}(M) \ := \ \{\langle X, x\rangle \mid X \in \mathcal{U} \cup \mathcal{V} \text{ and } x \in \mathsf{Val}(X)\} \ .$$

(For example, the shooting squad fluents $\dot{C}$ and $\neg C$ amount to $\langle C, 1\rangle$ and $\langle C, 0\rangle$.) Next we look at strings of subsets of $\mathsf{vv}(M)$. Given a string $s \in Pow(\mathsf{vv}(M))^*$, let $\mathsf{vv}(s)$ be the subset of $\mathsf{vv}(M)$ that appears in s — that is,

$$\mathsf{vv}(\alpha_1 \cdots \alpha_n) \ := \ \bigcup_{i=1}^{n} \alpha_i \ .$$

We say s is *M-compliant* if $\mathsf{vv}(s)$ is a function $f : (\mathcal{U} \cup \mathcal{V}) \to \bigcup_{X \in \mathcal{U}\cup\mathcal{V}} \mathsf{Val}(X)$ such that for all $X \in \mathcal{V}$,

$$f(X) \ = \ F_X(f_{-X})$$

where f_{-X} is f minus the variable-value pair $\langle X, f(X)\rangle$. Finally, let us call M *finite* if there are only finitely many variables, and every set $\mathsf{Val}(X)$ of values is finite (i.e., the union $\mathcal{U} \cup \mathcal{V} \cup \bigcup_{X \in \mathcal{U}\cup\mathcal{V}} \mathsf{Val}(X)$ is finite). Since $\mathsf{vv}(M)$ is finite for finite M, it follows that

Fact. *Given a finite causal model M, the set of M-compliant strings is a regular language.*

Note that for the shooting squad M, a string s is M-compliant iff $\mathsf{vv}(s)$ is either $\{\langle C, 0\rangle, \langle O, 0\rangle, \langle R, 0\rangle, \langle P, 0\rangle\}$ or $\{\langle C, 1\rangle, \langle O, 1\rangle, \langle R, 1\rangle, \langle P, 1\rangle\}$. The next sections provide a finite-state system that brings temporal order, paring the language of M-compliant strings down to

$$\boxed{\dot{C}}\boxed{\dot{O}}\boxed{\dot{R}}\boxed{\dot{P}} \ + \ \boxed{\neg C, \neg O, \neg R, \neg P}.$$

3 Finite-State Temporality

The relative conception of time as change and the associated notion of block reduction mentioned in the previous section can be understood through the example of a calendar year, represented by the string

$$s_{mo} \;:=\; \boxed{\text{Jan}}\boxed{\text{Feb}}\cdots\boxed{\text{Dec}}$$

of length 12 (with a month in each box), or (adding one of 31 days d1, d2,...d31) the string

$$s_{mo,dy} \;:=\; \boxed{\text{Jan,d1}}\boxed{\text{Jan,d2}}\cdots\boxed{\text{Jan,d31}}\boxed{\text{Feb,d1}}\cdots\boxed{\text{Dec,d31}}$$

of length 365 (a box per day in a non-leap year). Unlike the points in the real line $\mathbb{R}$, a box can split if we enlarge the set B of (*b*oxable) symbols we can put in it, as the change from $\boxed{\text{Jan}}$ in s_{mo} to $\boxed{\text{Jan,d1}}\boxed{\text{Jan,d2}}\cdots\boxed{\text{Jan,d31}}$ in $s_{mo,dy}$ illustrates. Or, reversing direction, from $s_{mo,dy}$ to s_{mo}, let us define two functions ρ_{mo} and bc that respectively, restricts B to the months $mo = \{\text{Jan,Feb,...Dec}\}$

$$\rho_{mo}(s_{mo,dy}) \;=\; \boxed{\text{Jan}}^{31}\boxed{\text{Feb}}^{28}\cdots\boxed{\text{Dec}}^{31}$$

and compresses a block α^n to α

$$\mathit{bc}(\boxed{\text{Jan}}^{31}\boxed{\text{Feb}}^{28}\cdots\boxed{\text{Dec}}^{31}) \;=\; \boxed{\text{Jan}}\boxed{\text{Feb}}\cdots\boxed{\text{Dec}} \;=\; s_{mo}$$

so that $\mathit{bc}(\rho_{mo}(s_{mo,dy})) = s_{mo}$. More precisely, for $A \subseteq B$, ρ_A sees only the elements of A (discarding non-A's)

$$\rho_A(\alpha_1\alpha_2\cdots\alpha_n) \;:=\; (\alpha_1 \cap A)(\alpha_2 \cap A)\cdots(\alpha_n \cap A)$$

whereas *block compression* bc sees only change (discarding repetitions/stuttering)

$$\mathit{bc}(s) \;:=\; \begin{cases} \mathit{bc}(\alpha s') & \text{if } s = \alpha\alpha s' \\ \alpha\,\mathit{bc}(\alpha' s') & \text{if } s = \alpha\alpha' s' \text{ with } \alpha \neq \alpha' \\ s & \text{otherwise.} \end{cases}$$

Let bc_A be the composition $\rho_A; \mathit{bc}$ mapping s to

$$\mathit{bc}_A(s) \;:=\; \mathit{bc}(\rho_A(s)),$$

so that

$$\mathit{bc}_{\{\text{Jan}\}}(s_{mo,dy}) = \mathit{bc}_{\{\text{Jan}\}}(s_{mo}) = \boxed{\text{Jan}}\boxed{} \qquad \mathit{bc}_{\{\text{d1}\}}(s_{mo,dy}) = (\boxed{\text{d1}}\boxed{})^{12}$$

$$\mathit{bc}_{\{\text{Feb}\}}(s_{mo,dy}) = \mathit{bc}_{\{\text{Feb}\}}(s_{mo}) = \boxed{}\boxed{\text{Feb}}\boxed{} \qquad \mathit{bc}_{\{\text{d2}\}}(s_{mo,dy}) = (\boxed{}\boxed{\text{d2}})^{12}\boxed{}\,.$$

We can delete any initial or final empty boxes by a function *unpad*, which we apply after bc_A to form π_A

$$\pi_A(s) \;:=\; \mathit{unpad}(\mathit{bc}_A(s)).$$

A symbol $e \in B$ is then defined to be an *s-interval* if $\pi_{\{e\}}(s)$ is $\boxed{e}$

$$s \models \text{interval}(e) \quad :\Leftrightarrow \quad \pi_{\{e\}}(s) = \boxed{e}\ .$$

For relations between intervals, we apply $\pi_{\{e,e'\}}$ to the set of strings s where e and e' are s-intervals to form

$$\mathcal{L}_\pi(\{e,e'\}) \quad := \quad \{\pi_{\{e,e'\}}(s) \mid \pi_{\{e\}}(s) = \boxed{e} \text{ and } \pi_{\{e'\}}(s) = \boxed{e'}\}.$$

There are 13 strings in $\mathcal{L}_\pi(\{e,e'\})$, one per interval relation in Allen 1983, refining the relations $\prec$ of *(complete) precedence* and $\bigcirc$ of *overlap* used in the Russell-Wiener construction of time from events; see Table 1. We have

Table 1. From Russell-Wiener to Allen

RW	**Allen**	$Pow(\{e,e'\})^*$	**Allen**	$Pow(\{e,e'\})^*$	**Allen**	$Pow(\{e,e'\})^*$
$e \bigcirc e'$	$e\ \texttt{=}\ e'$	$\boxed{e,e'}$	$e\ \texttt{fi}\ e'$	$\boxed{e}\boxed{e,e'}$	$e\ \texttt{f}\ e'$	$\boxed{e'}\boxed{e,e'}$
	$e\ \texttt{si}\ e'$	$\boxed{e,e'}\boxed{e}$	$e\ \texttt{di}\ e'$	$\boxed{e}\boxed{e,e'}\boxed{e}$	$e\ \texttt{oi}\ e'$	$\boxed{e'}\boxed{e,e'}\boxed{e}$
	$e\ \texttt{s}\ e'$	$\boxed{e,e'}\boxed{e'}$	$e\ \texttt{o}\ e'$	$\boxed{e}\boxed{e,e'}\boxed{e'}$	$e\ \texttt{d}\ e'$	$\boxed{e'}\boxed{e,e'}\boxed{e'}$
$e \prec e'$	$e\ \texttt{m}\ e'$	$\boxed{e}\boxed{e'}$	$e < e'$	$\boxed{e}\boxed{\ }\boxed{e'}$		
$e' \prec e$	$e\ \texttt{mi}\ e'$	$\boxed{e'}\boxed{e}$	$e > e'$	$\boxed{e'}\boxed{\ }\boxed{e}$		

$$\mathcal{L}_\pi(\{e,e'\}) \quad = \quad \text{Allen}(e \bigcirc e') + \text{Allen}(e \prec e') + \text{Allen}(e' \prec e)$$

where Allen$(e \bigcirc e')$ consists of the 9 strings in which e overlaps e'

$$\text{Allen}(e \bigcirc e') \quad := \quad (\boxed{e} + \boxed{e'} + \epsilon)\boxed{e,e'}(\boxed{e} + \boxed{e'} + \epsilon)$$

(with empty string ϵ), and Allen$(e \prec e')$ consists of the 2 strings in which e precedes e'

$$\text{Allen}(e \prec e') \quad := \quad \boxed{e}\boxed{e'} + \boxed{e}\boxed{\ }\boxed{e'}$$

and similarly for Allen$(e' \prec e)$. If e and e' are s-intervals, we have

$$s \models e\ R\ e' \ \Leftrightarrow\ \pi_{\{e,e'\}}(s) \in \text{Allen}(e\ R\ e') \qquad \text{for } R \in \{\bigcirc, \prec, \succ\}$$

or, for example, in the case of the Allen relation `f` (for "*f*inish"),

$$s \models e\ \texttt{f}e' \quad \Leftrightarrow \quad \pi_{\{e,e'\}}(s) = \boxed{e'}\boxed{e,e'}.$$

Returning to the previous section, we can strengthen the notion of an M-compliant string s by requiring that for all $\langle X, x\rangle, \langle Y, y\rangle \in \mathsf{vv}(s)$ such that X depends on Y, $\pi_{\{\langle X,x\rangle,\langle Y,y\rangle\}}(s)$ is one of the strings corresponding, under Table 1, to the desired relation of precedence. (That relation may say, be weakened from $\prec$ to include the Allen subrelations `=,si,s,fi,di,o` of Russell-Wiener overlap $\bigcirc$.)

To check if a string satisfies a formula φ, the approach taken above is to choose a finite set A of symbols such that the function π_A picks out a suitable granularity at which to analyze φ. The computational pressure is to make A as small as possible, but other formulas may require finer granularities, necessitating an enlargement of A. For this reason, it may be useful to consider an infinite set E of symbols and its family $Fin(E)$ of finite subsets, forming $Fin(E)$-indexed strings $(s_A)_{A\in Fin(E)}$ in which s_A can be calculated as $\pi_A(s_B)$ for any $B \supseteq A$, and collecting these in the *inverse limit of* the system $(\pi_A)_{A\in Fin(E)}$ of functions

$$\varprojlim (\pi_A)_{A\in Fin(E)} := \{(s_A)_{A\in Fin(E)} \in \prod_{A\in Fin(E)} Pow(A)^* \mid \\ s_A = \pi_A(s_B) \text{ whenever } A \subseteq B \in Fin(E)\} .$$

As a function between strings over the alphabet $Pow(B)$, π_A is computable by a finite-state transducer (for finite B). Hence, if the π_A-based approach just outlined works for a formula φ, the set of strings satisfying φ is regular — making φ equivalent to a formula of Monadic Second-Order logic (e.g. Thomas 1997). While finite-state methods need not proceed via π_A, what the inverse limit above offers is a system of bounded but refinable temporal granularities in which intervals are conceptually prior to points (as a box splits into a string of boxes, upon closer examination). More in Fernando 2011a.

4 Inertia and Bounded Entailments

An assumption implicit in the previous section, for instance Table 1, is that a string such as $\boxed{e}\boxed{}\boxed{e'}$ can be read as $\boxed{e, \neg e'}\boxed{\neg e, \neg e'}\boxed{\neg e, e'}$, and more generally, $\alpha_1 \cdots \alpha_n$ is equivalent to $cl_\neg(\alpha_1 \cdots \alpha_n) := \alpha'_1 \cdots \alpha'_n$ where

$$\alpha'_i := \alpha_i \cup \{\neg e \mid e \in \bigcup_{j=1}^{n} \alpha_j - \alpha_i\} \quad \text{for } 1 \leq i \leq n.$$

As information tends to get conveyed partially, it is useful to accommodate underspecification in a string by retracting this assumption (replacing strings s in Table 1 by $cl_\neg(s)$). Rather than forming $cl_\neg(s)$ freely, we insert $\neg e$ whenever appropriate (as we do when replacing strings s in Table 1 by $cl_\neg(s)$).[4] Furthermore, we can treat inertial fluents e as follows, in accordance with condition (I) from section 2. We suppose there are forces acting for and against e, and fluents fe decreeing "there is a force for e" — at least, that is, in strings such that whenever fe occurs, e holds at the next moment unless some force opposes it. More precisely, let $\mathcal{F}_e$ be the set of strings $\alpha_1 \cdots \alpha_n$ such that

$$\text{f}e \in \alpha_i \quad \text{implies} \quad e \in \alpha_{i+1} \text{ or } \text{f}\neg e \in \alpha_i \qquad \text{for } 1 \leq i < n$$

[4] There is a distinction here between a string qua index versus a string qua denotation drawn at length in Fernando 2011. In the former case, it is natural to equate s with $cl_\neg(s)$; not so in the latter. Note $cl_\neg$ is computable by a finite-state transducer.

(where f$\neg e$ effectively says "there is a force opposed to e"). $\mathcal{F}_e$ can be expressed succinctly as

$$\boxed{\mathrm{f}e}\boxed{} \Rightarrow \boxed{}\boxed{e} + \boxed{\mathrm{f}\neg e}\boxed{}$$

using a binary operation $\Rightarrow$ that maps a pair of regular languages to a regular language (Fernando 2008, 2011). Closely related to $\mathcal{F}_e$ is the regular language

$$\boxed{e}\boxed{} \Rightarrow \boxed{}\boxed{e} + \boxed{\mathrm{f}\neg e}\boxed{}$$

which intersected with

$$\boxed{}\boxed{e} \Rightarrow \boxed{e}\boxed{} + \boxed{\mathrm{f}e}\boxed{}$$

yields a regular language $\mathcal{I}_e$ in which e persists forward and backward in the absence of forces on it. It is natural to view the languages $\mathcal{F}_e$ and $\mathcal{I}_e$ as constraints (satisfied by the strings belonging to them), associated with states (as opposed to events) represented by e.

An example of a state that an inertial fluent may represent is given by the habitual (3a), repeated below, followed by the episodic (3b).

(3) a. Tess eats dal.
 b. Tess is eating dal.

We must be careful not to confuse a string s representing an instance of Tess eating dal with the fluent e representing the habit of Tess eating dal. We may, however, expect the string s to belong to a language $\mathcal{L}_e$ associated with e. That is, $\mathcal{L}_e$ sits alongside the sets $\mathcal{F}_e$ and $\mathcal{I}_e$ above as languages associated (in different ways) with e. The obvious question is what is the status of such languages in our semantic theory?

I close with the suggestion that such languages are specifications of causal structures required by the *rules-and-regulations* view of Carlson 1995 to ground the truth of generic sentences. (These languages are underspecifications insofar as the causal structures are automata, but at present, I see no reason for that additional specificity.) Beyond $\mathcal{F}_e, \mathcal{I}_e$ and $\mathcal{L}_e$, there are *episodes* in the sense of Moens and Steedman 1988 consisting of "sequences of causally or otherwise contingently related sequences of events" that (lest we confuse these with episodic instances at the world) are "more related to the notion of a plan of action or an explanation of an event's occurrence than to anything to do with time itself" (page 26). The intuition is that these languages are

> resources for constructing local languages for *use* in particular situations

to quote Cooper and Ranta 2008 slightly out of context. A shift is made here from monolithic truth to use (or action) that is very much in line with the focus in ST_{fa} on finite-state transducers (rather than say, the real line $\mathbb{R}$ or possible worlds). More questions are certainly left open than answered, an important issue being inference. For entailments between formulas φ and ψ based on satisfaction $\models$, we can use a language L to relativize the inclusion

$$\varphi \vdash_L \psi \;:\Leftrightarrow\; (\forall s \in L)\; s \models \varphi \text{ implies } s \models \psi$$

which is decidable provided φ and ψ are in Monadic Second Order Logic and L is regular. In fact, the inclusion remains decidable for L context-free, as observed by Makoto Kanazawa. Regular or not, L may consist of strings that represent episodic instances beyond those of any single world. It is one thing to say a causal force for L exists at a world and time, another to spell out the consequences for the episodic instances at a world and time, and do justice to competing "inference tickets" (Ryle 1949).

References

Allen, J.F.: Maintaining knowledge about temporal intervals. Communications of the ACM 26(11), 832–843 (1983)

Bird, A.: Dispositional Expressions. In: The Routledge Companion to Philosophy of Language (to appear)

Carlson, G.N.: Truth conditions of generic sentences: two contrasting views. In: The Generic Book, pp. 224–237. University of Chicago Press (1995)

Cooper, R., Ranta, A.: Natural languages as collections of resources. In: Language in Flux: Dialogue Coordination, Language Variation, Change and Evolution. College Publications, London (2008)

Dowty, D.: Word Meaning and Montague Grammar. Reidel (1979)

Fernando, T.: A finite-state approach to events in natural language semantics. J. Logic and Computation 14(1), 79–92 (2004)

Fernando, T.: Branching from inertia worlds. J. Semantics 25(3), 321–344 (2008)

Fernando, T.: Regular relations for temporal propositions. Natural Language Engineering 17(2), 163–184 (2011)

Fernando, T.: Finite-state representations embodying temporal relations. In: 9th International Workshop on Finite State Methods and Natural Language Processing, Blois, pp. 12–20 (2011a), A revised, extended version is in the author's webpage

Halpern, J.Y., Pearl, J.: Causes and explanations: a structural-model approach, Part 1: Causes. Brit. J. Phil. Sci. 56, 843–887 (2005)

Hiddleston, E.: A causal theory of counterfactuals. Noûs 39(4), 632–657 (2005)

Lück, U.: Continu'ous Time Goes by Russell. Notre Dame J. Formal Logic 47(3), 397–434 (2006)

Moens, M., Steedman, M.: Temporal ontology and temporal reference. Computational Linguistics 14(2), 15–28 (1988)

Pearl, J.: Causality: Models, Reasoning and Inference, 2nd edn. Cambridge University Press (2009)

Ryle, G.: The Concept of Mind. University of Chicago Press (1949)

Schulz, K.: If you'd wiggled A, then B would've changed Causality and counterfactual conditionals. Synthese 179(2), 239–251 (2011)

Steedman, M.: The Productions of Time. Draft tutorial notes about temporal semantics, Draft 5.0 (May 2005), `http://homepages.inf.ed.ac.uk/steedman/papers.html`

Thomas, W.: Languages, automata and logic. In: Handbook of Formal Languages: Beyond Words, vol. 3, pp. 389–455. Springer (1997)