How Telicity Creates Time

Östen Dahl

Abstract: Most treatments of temporal semantics start out from the conception of time as a line stretching from the past into the future, which is then populated with eventualities or situations. This paper explores how time can be seen as emerging from the construction of representations of reality in which the basic building blocks are static—i.e., timeless—representations, which are connected to each other by events that are transitions between them and that create an ordering which can be understood as temporal. This connects to von Wright's "logic of change" and the "hybrid semantics" suggested by Herweg and Löbner. In this context, telicity is seen as the capacity of events, or of the predicates that express them, to "create time" in the sense of defining a before and an after. The basic elements of the model are global states, which are timeless taken in isolation but are connected by transition events, which transform one global state into another and thereby define the temporal relationships between them. Transition events, corresponding to Vendlerian achievements, represent simple changes which are then the basis for all other constructs in the model, most notably delimited states, Vendlerian activities (atelic dynamic eventualities), and accomplishments (telic non-punctual eventualities), but also time points and intervals. Transition events are further instrumental in constructing narrative structures and are responsible for narrative progression.

1. Introduction: The Nature of Time

Tense and aspect systems are notoriously difficult to get a grip on, in particular as concerns their semantics. It is true that the underlying conceptual categories are particularly elusive but it is also the case that we tend to take many things about them for granted. Therefore, even if it cannot be the task of linguists to fully elucidate the reality behind categories such as time, it is sometimes useful to question the ways in which we use these categories in our theorizing about language.

One central issue is how time is related to change. In *Physics*, Aristotle raises the question, "Is time independent of change?" and argues that the answer is no, "for when the state of our own minds does not change at all, or we have not noticed its changing, we do not real-

Journal of Slavic Linguistics 21(1): 45–76, 2013.

ize that time has elapsed". It would seem that this argument does not so much concern the ultimate nature of time as our way of understanding it. On the other hand, the latter question is also what primarily concerns us as linguists, as we want to understand the conceptual underpinnings of what we express in language and what is reflected in the grammars and lexicons of human languages. But since we are biological organisms which are part of the physical world, this striving also leads us to investigate how our cognitive apparatus helps us to survive in that world. We thus want to understand what our cognitive representations of reality are like, but also how these representations map onto reality. Furthermore, when speaking of time, cognition, and language, it is important to remember that speaking and thinking are in themselves processes that take place in time and we are creatures that live in time. What we need is thus what I have elsewhere (Dahl 2007) called an ecological semantics based on a cognitively sound ontology**.** In that context, Aristotle's argument makes sense: change is indeed essential for our experience of time.

Most treatments of temporal semantics start out from the conception of time as a line stretching from the past into the future, which is then populated with eventualities or situations, which are not entirely happy cover terms for states, processes, and events. However, the timeline as a graphic representation is a relatively new invention, which appeared only in the eighteenth century in Western culture (Rosenberg and Grafton 2010). Discussing how we date events in memory, the psychologist William Friedman uses the expression "the chronological illusion" (Friedman 1993) to characterize the idea of time being mentally represented as a "unitary, linear continuum assigning each event a unique location" (44). Instead, our view of the past "depends on a process of active, repeated construction" involving "coincidence, locations in recurrent patterns, and independent sequences of meaningfully related events" (62).

In this paper I want to explore how time can be seen as emerging from the construction of representations of reality in which the basic building blocks are static—i.e., timeless—representations, which are connected to each other by events[1] that are transitions between them

[1] The term event is a bit dangerous. Some scholars use it as a cover term equivalent to eventuality or situation. I will use it more restrictively, in a way that I think is more in accordance with everyday usage. Most importantly, events always involve change.

and create an ordering which can be understood as temporal. The most obvious kind of linguistic manifestation of such a construction process are narratives. The traditional concept of narrative progression or moving the narration forward can be understood as a stepwise construction of a temporal representation of reality usually involving both static representations and events. The primary means for introducing "time-creating" events in a narrative are telic predicates. I am using the term telic in a broad sense in which a predicate is telic if it involves a specific limit to the eventuality spoken of (see further section 5 below).

In a language such as Russian, verbs in telic predicates in this way will normally be in the perfective aspect, but the phenomenon of narrative progression is universal and goes beyond particular grammatical systems.

A few notes on ontology are in order here. In this paper I will often speak of reification and of constructs. On a naïve view, the world consists of things or objects with an independent existence. An alternative is to think of the things we speak about as constructed out of patterns that we discern in the world around us. These may not correspond to things in the ordinary, physical sense. For instance, think of a man who is digging a hole in the ground, standing in the sun. Neither the hole nor the shadow that we see behind the man is a physical object; yet we can easily think of them as such. Moreover, they arguably have an objective existence in the sense that they are visible and their location can be objectively determined. By the process of reification we create objects that we can speak about; we treat them as entities which exist, can be referred to, quantified over, and have properties that can be predicated of them. But these properties can in their turn be reified. If I say *Mary is kind*, I predicate something of Mary; but if I say *Mary's kindness is surprising* I predicate something not of Mary but of her kindness, which is now treated as a reified entity in its own right. Reification of abstract entities is a general precondition for processes that create nouns and noun phrases, as discussed in Langacker 2008.

To understand the nature of reification and constructs, consider a notion which is closely related to that of telicity, namely that of a

Thus, *Peter woke up* describes an event but *Peter is awake* does not. I have realized that some of the difficulties experienced by readers of the draft version of this paper depended on my not having made this clear enough.

boundary. Boundaries are ubiquitous in daily life and in our thinking. In mathematics, more specifically in topology, boundaries are usually thought of as sets of points in space, as in the Wikipedia definition: "a boundary of a subset S of a topological space X is the set of points which can be approached both from S and from the outside of S". If we are talking of a segment of a line, the right boundary of that segment would be either the rightmost point of that line or the leftmost point not belonging to the segment.

Linguists tend to think of boundaries in another way, however. A compound word such as *blackbird* consists of two morphemes, *black* and *bird.* In linguistics, we see these morphemes as separated by a boundary, but we do not think of such a boundary as either the last segment of the first morpheme or the first segment of the second. Rather, we think of the boundary as going between the morphemes, and thus in a way being an entity which is both different in nature and separate from both of them, and which can be marked by a special symbol, e.g., a hyphen: *black-bird*. Likewise, in standard orthography, the boundaries between words are marked by spaces which are not part of words themselves. In fact, this is the natural way of treating boundaries in discrete systems. It means that boundaries between elements in the system are a different kind of animal than the elements themselves, which may seem a bad thing from the point of view of ontological parsimony, but is more in line with how we think of boundaries in everyday life. This does not mean that boundaries are independent of the elements of the system—they are normally defined by the properties of the elements they separate, as in the picture below, where we have a boundary between the white and black squares:

(1)

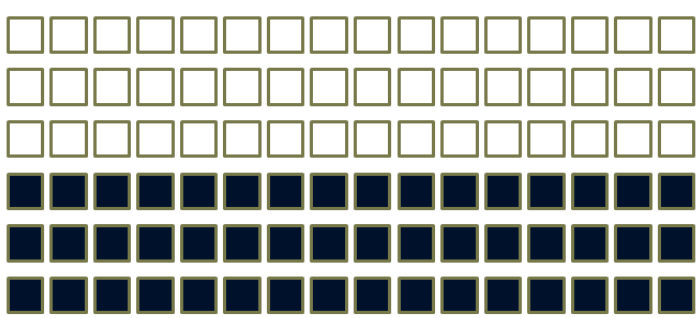

Boundaries in a system are thus constructs based on the properties of elements of the system.

As noted above, theories of temporal reference tend to model time as a line (or technically a vector), and temporal entities such as events or states are characterized by their location on such a line. Such a timeline is a continuous representation of a dimension of reality, similar to a spatial representation such as a two-dimensional map. But representations need not be continuous, they can also be discrete**.** Even some maps are better understood as discrete, for instance, subway (underground) maps of the classical kind. What does a subway map show? Basically, it displays a structure consisting of a set of individual objects (stations) and connections between these objects. What makes the map a discrete representation rather than a continuous one is that the distance between the stations is not shown in a systematic way: a line between two stations tells us that the stations are connected but not how far they are from each other. Of course, the stations in an underground map can be assigned locations on the topographic (continuous) map, and eventually also in reality. But if we just want to know how to get from one point to another, this is something that we need not worry about. In a similar way, I may understand and enjoy a story that I read without knowing how to locate it in space and time (or even without knowing whether it is true or not).

In fact, discrete representations or models of time show up in many contexts. One obvious type of case is found in games such as chess or tic-tac-toe which are played in discrete moves. For the course of such a game, the length of the physical time intervals is irrelevant,[2] and would normally not be represented in an account of it. Another type of situation when discrete representations are used is when observations of some aspect of reality are made at regular (or irregular) intervals. If we measure the weight of a baby every morning, we'll get a series of observations which in effect constitute a discrete model of the baby's development. Statements that we make about it such as "The baby is steadily gaining weight" will be based only on the daily

[2] Chess games, as played in tournaments, have a more complex time structure, since there is a restriction on the total time a player can use for his moves, and this is also recorded. In fact, it could be argued that there are three layers of time in such a chess game: one discrete level, measured purely in moves, and two continuous ones, normal physical time and game time measured by the chess clocks (which would be different from physical time in not including breaks).

observations, and any fluctuations in between will be neglected. In the world of computers, discrete time systems are prominent. Cassandras and Lafortune (2008: 25) enumerate a number of reasons for the importance of such systems, the most obvious being that the internal clocks of computers make them operate in discrete time fashion. "Discrete event simulation" represents a system as a sequence of events, where the events are instantaneous changes or transitions between states in the system.

2. From States to Events

I shall now try to show how we can build the ontology we need for a temporal semantics in a stepwise fashion, starting out from representations of the world that do not involve time.

As an illustrative example, consider a world or system consisting of three elements or objects denoted by the letters a, b, and c. Each object has a value represented by an integer. We can represent the system as follows:

(2)

a=2
b=3
c=4

We make the assumption that the elements are constant in the system, their values are not fixed but can vary between different states of the system as in (3).

(3)

a=2	a=4	a=4	a=4
b=3	b=3	b=2	b=2
c=4	c=4	c=4	c=6

The set of all possible states makes up the state space of the system. This does not imply any temporal relationship between the states or that the states hold at some point in time. So far they exist only as logical possibilities (in a similar way to the set of sentences generated by a phrase-structure grammar).

Before going further, I should point out that the term "state" is multiply ambiguous. As used in the previous paragraph, "state" refers to the totality of what can be said of an object or a system—often at a specific point in time, but also just with reference to a set of possibilities, without the assumption of any temporal dimension, as in our example. We can call this a global state. The word "state" is often used in the sense of "global state", as exemplified by expressions such as "the State of the Union address". In the temporal semantics literature, however, we more frequently find the word "state" associated with the content of predicates such as Russian *byt' sčastlivym* or English *be happy*, which single out one particular aspect of the global state of an individual. There is a further ambiguity here, though, that tends to be glossed over. If Smith is drunk on Monday and then again on Wednesday, we may say that he was in the same state twice. But we may also say that his first drunken state lasted longer than his second one. In the first case, we are speaking of a generic state, in the second of a specific state. Authors tend to switch happily between the two interpretations, but the difference between them is of a certain importance, as I shall argue below.

Strictly speaking, what we see in a representation such as (3) is not a set of states—in the sense of global states—but rather what we can call state descriptions. The distinction is motivated since there is no necessity for a state description to be total. In other words, a state description is a partial or total characterization of a global state.

An important step towards including a temporal dimension in (3) is to introduce operations by which one state is transformed into another. We can for instance derive the second state description in (3) by adding 2 to the value of the element a. Using programming language notation, we can represent this as follows:

(4)

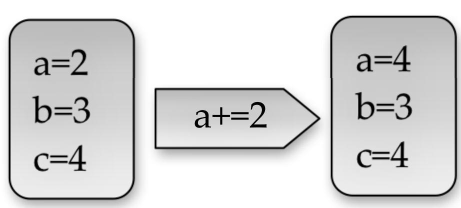

In principle, we may still think of this as a representation of the logical relationship between two possible states, but we may also see the relationship as a temporal one: we first have the state to the left, then it is transformed into the state to the right. In this interpretation, the transformation of the first state into the second can be seen as a transition that takes place in physical or imagined time, or in simpler terms, as a kind of event, that we can call a transition event. There is a duality relationship between states and transition events, in that, on the one hand, the source state (the state to the left in (4)) and the transition event together unequivocally determine the target state (the state to the right in (4)) and, on the other, the two states together unequivocally determine the transition. It follows that a representation such as (3), interpreted as a temporally ordered sequence of states, and a representation such as (5), containing a specification of an initial state and a sequence of events, are equivalent, that is, contain the same information:

(5) *Initial state*

a = 3
b = 3
c = 2

Transition events

a = +2
b = +1
c = –1

A similar equivalence can be found in two-dimensional representations of the world such as drawings and maps. An area can be shown either as a blob of color or as a contour, that is, displaying its boundary as a line. Thus, the following representations are equivalent:

(6)

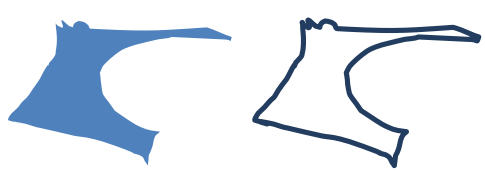

Since a representation in terms of events will in general be considerably more economical than one in terms of state descriptions, it is also preferable in most situations. It would for instance be possible to describe a chess game by describing the successive states of the game (the boards), but usually it's done by enumerating the transformations, i.e., the moves. Likewise, narratives often contain a description of the initial setting followed by a series of events told in the order they happened. The illustration in (7) from Lewis Carroll's *Through the Looking Glass* is a combination of both these possibilities.

(7)

RED.

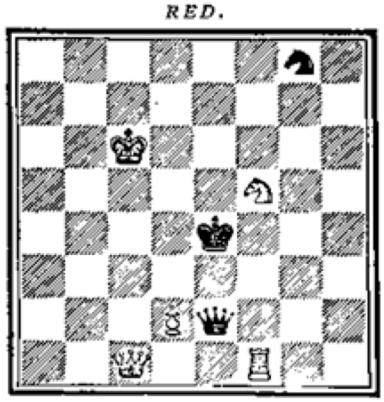

WHITE.

White Pawn (Alice) to play, and win in eleven moves.

	PAGE		PAGE
1. Alice meets R. Q. . .	35	1. R. Q. to K. R's 4th.	41
2. Alice through Q's 3d (*by railway*)	44	2. W. Q. to Q. B's 4th (*after shawl*)	75
to Q's 4th (*Tweedledum and Tweedledee*)	49		
3. Alice meets W. Q. (*with shawl*)	75	3. W. Q. to Q. B's 5th (*becomes sheep*) . . .	81
4. Alice to Q's 5th (*shop, river, shop*)	82	4. W. Q. to K. B's 8th (*leaves egg on shelf*)	90
		5. W. Q. to Q. B's 8th (*flying from R. Kt.*)	116
5. Alice to Q's 6th (*Humpty Dumpty*)	90	6. R. Kt. to K's 2d (check)	122
6. Alice to Q's 7th (*forest*)	120	7. W. Kt. to K. B's 5th	140
7. W. Kt. takes R. Kt.	125	8. R. Q. to K's sq. (*examination*)	143
8. Alice to Q's 8th (*coronation*)	141	9. Queens castle	156
9. Alice becomes Queen	142	10. W. Q. to Q. R's 6th (*soup*)	162
10. Alice castles (*feast*)	156		
11. Alice takes R. Q. and wins	163		

[10]

From sequences of abstract structures, let us now proceed to how we perceive the world in physical time. A picture is a static representation of the world. Consider, however, the two pictures in (8) and suppose that they are projected on a screen in rapid succession.

(8)

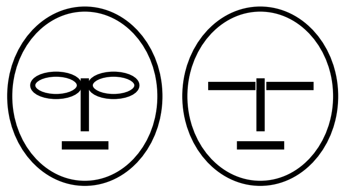

Someone who watches the screen will see this not as two separate pictures but rather as one and the same picture undergoing a change. Moreover, we feel that we have witnessed an event that we can talk about using a dynamic verb, such as *zakryt'* in Russian or *close* in English:

(9) Petr zakryl glaza.

(10) Peter closed his eyes.

Von Wright (1963) in his "logic of change" defines an operator *T* which can be read as "and then". A formula *pTq* thus means '*p* and then *q*', that is, there is first a state of affairs in which *p* is true and then a state of affairs where *q* is true. But *pTq* according to von Wright describes an event which consists in the change or transition from one state to another.

Do events really exist? If we look at (8) we see two static views of an object. There is nothing we can identify with an event. We can only speak of an event when we see the two pictures as two successive states of that object and identify the way the object changes between those two states. Von Wright says that such an event can be seen as "an ordered pair of two states of affairs" but also that "the event 'itself' is the change or transition" from the first state to the second. The first formulation can be seen as expressing a reductionist view, on which the event is nothing but the two ordered states. The second formulation, on the other hand, suggests an emergentist view on which the event is seen as something in its own right, where it is reified, made into a thing, in the sense discussed above. Philosophers have had different opinions about the necessity or desirability of reifying events, with Donald Davidson as the foremost proponent of what I here call the emergentist view (Davidson 1967) and Quine as the eternal defender of a parsimonious ontology (Quine 1985).

From a cognitive point of view, however, there seems to be little doubt that we think of events as entities in themselves, and this is also reflected in the way we speak about them. Actually, the very fact that we can speak about events implies that they are seen as entities, albeit abstract ones. Arguably, events are constructs, in the sense defined above. Fleischman (1990: 99) articulates the view that

> the event is but a hermeneutic construct for converting an undifferentiated continuum of the raw data of experience, or of the imagination, into the verbal structures we use to talk about experience: narratives, stories… a construct that mediates between experience and language, yet belongs strictly to neither domain.

I do not want to go as far as Fleischman seems to do when she says "events only come into being as a result of linguistic operations" (1990: 98). Rather, the reification of events is a way of making sense of what she calls "the raw data of experience", a process that certainly takes place even in the absence of language. Note that watching the two images in (8) being projected on the screen in rapid succession, we cannot help perceiving it as an event. The reification here takes place at a subconscious level and is independent of whether we express it in language or not.

3. The Detector Model

To see if it is possible to grasp how the reification of events can come about, I shall now introduce a maximally simple information processing device that I shall refer to as a detector. There is nothing particularly original in this notion; analogues to detectors are found under many different names in the literature. Basically, a detector is a device that accepts input from the environment and outputs signals on the basis of the input. A simple example of a detector is a light (e.g., on the dashboard of a car) that turns on when a certain temperature (say 20 degrees) is reached. The detectors that I shall describe here react to input in accordance with stimulus-response rules (Holland 1995: 7) of the form

(11) IF condition$_1$ THEN signal$_1$ (ELSE signal$_2$)

Suppose for instance that we have a device, similar to the dashboard light, that receives input in the form of digits (which we may interpret as giving information about the temperature in the environment) and outputs binary signals according to some rule such as the following:

(12) If input > 20 then output "1" else output "0"

The following picture illustrates the workings of such a device:

(13)

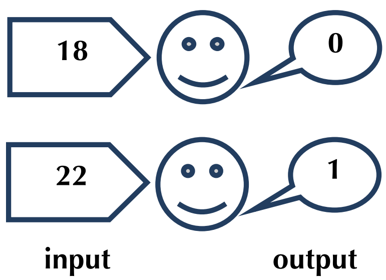

What we can note is that the detector outputs maximally simple state descriptions: '1' meaning 'the temperature is above 20°, and '0', meaning 'the temperature is not above 20°. To be able to process also information about events, the detector crucially needs to be able to store previously received input. In the maximally simple case, a detector which is able to "remember" the immediate preceding item will be able to restrict its output to situations where a change has taken place, by comparing the input with the information it has stored. Suppose for instance that the detector follows a rule such as

(14) If input > 20 and store < 20 then output "1" else do nothing

The following diagram illustrates the application of (14):

(15)

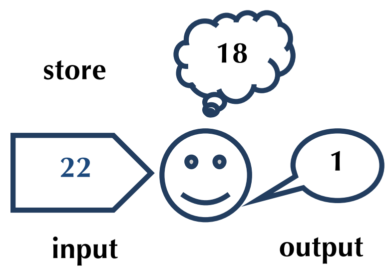

The detector described so far has a rather limited cognitive capacity: it can store information about one state of affairs and compare it to incoming information. (Compare a telephone that stores the last number dialed.) It may be noted here that comparing a stored representation of the world with incoming perceptual information is a very basic cognitive operation in living organisms, something that we do constantly. It does not presuppose an understanding of time, however. Arguably, the stored information is in itself atemporal, since it contains one state of affairs only. It might be said that the detector itself has no conception of time. Moreover, there is no necessity to interpret the relationship between the stored and the incoming information in temporal terms. Suppose that I believe that the temperature is 18 degrees and I hear someone utter the Russian sentence (16) or the English sentence (17).

(16) Temperatura—22 gradusa.

(17) The temperature is 22 degrees.

If (17) is true, there are two logical possibilities: either the temperature has changed or I was wrong and it has been 22 degrees all along. In many situations, the choice between these two alternatives may not be of any great significance. However, a Russian sentence such as (18) or an English sentence such as (19) unequivocally signals that an event involving a change in temperature has taken place.

(18) Temperatura podnjalas' do 22 gradusov.

(19) The temperature has risen to 22 degrees.

Examples (18) and (19) are most naturally used in a situation where a lower temperature has been registered on a previous occasion, that is, where we already have a static representation of the world. It can be seen as an invitation to expand that representation into a sequence of two temporally related states of affairs, related to each other by an event in such a way that one holds before and the other after the event. Examples (18) and (19) thus **create time** in the metaphorical sense of introducing temporal structure in a representation of the world. For a linguistic expression to create time in this way, it has to introduce an event, either through its conventional meaning or through some kind of invited inference or implicature. Examples (16–17) and (18–19) clearly differ in this respect.

4. Hybrid Semantics

Let us now see what kind of ontology the preceding discussion has resulted in. In the first step we introduced representations of global states, which can be said to be the basic elements of temporal structures. In the second step transition events which transformed one such state into another were added. Global states and transition events are then the basic elements of temporal structures. Temporal ordering relations are derived from the order in which transition events are applied to states, and points in time are best seen as constructs which are abstract positions in a temporal structure, comparable to the positions on a waiting list. We can speak of the second position on the list, abstracting away from what individual has that position.

This ontology is in keeping with the idea of a hybrid semantics (Löbner 1988, Herweg 1991). Such a semantics combines elements of two types of semantic theories: traditional propositional or truth-conditional semantics and eventuality semantics. In propositional semantics a predicate is technically nothing but a function from sequences of individuals to truth values. The developments within formal semantics in the Montague Semantics tradition have tended to be based on this kind of framework. By contrast in eventuality semantics, both static and dynamic predicates are generally assumed to introduce entities variously referred to as situations or eventualities. In principle then, every predication involves the reification of a temporal entity. Examples of treatments adhering to eventuality semantics are found in particular within the tradition of Discourse Representation Theory (for a recent treatment, see Kamp, Genabith, and Reyle 2011), but the general ideas are quite widespread in works on temporal semantics. Most commonly, treatments follow Vendler (1957) in assuming a taxonomy consisting of states, activities (dynamic, atelic, non-punctual), accomplishments (dynamic, telic, non-punctual), and achievements (dynamic, telic, punctual). For instance, (20–21) are assumed to introduce a state, a stative situation/eventuality which is linked to Peter and the time period during which he is happy.

(20) Petr sčastliv.

(21) Peter is happy.

Similarly, (22–23) introduce an achievement, a telic, punctual, and dynamic situation/eventuality linked to Peter and the moment when he became happy.

(22) Petr stal sčastlivym.

(23) Peter became happy.

I noted above that the terms "eventualities" and "situations" are not entirely felicitous as cover terms for the entities covered by the Vendlerian taxonomy. It is rather striking that languages tend to lack an everyday word that could be used for this purported general concept. One possible reason is that eventualities do not, after all, form a natural class. This is in fact what is implied by the way they are treated in hybrid semantics, where reification is more restricted than in eventuality semantics and applies primarily to predications involving events. Events are seen as temporal individuals,[3] that is, they are individual things with temporal relations to each other, while stative predicates are just that, predicates which characterize global states. There are good reasons to think of eventive and stative predicates as not being on the same level. They have rather different logical properties and also tend to be treated quite differently in the grammars of human languages, in particular in tense and aspect systems (Dahl

[3] "Individual" is here used in the sense common in philosophy and logic, simply meaning "individual object/entity" and not implying personhood.

1985: 28). The assumption made in eventuality semantics that stative predicates like those in (20–21) introduce stative eventualities leads in effect to assimilating states to individual entities such as events, as noted by Herweg (1991: 373).

The reason for coming up with the idea of eventualities as a general ontological category seems to have been that scholars have tried to find something that can serve as the semantic correlate of any predicate—or at least of any verb. In fact, the discussion has often been restricted to those, as when Vendler's seminal paper was called "Verbs and Times" with non-verbal predication passed over in silence. Indeed, it is not obvious what kind of eventuality would underlie a sentence such as *Two plus two equals four*. In this connection, the different versions of Carlota Smith's theory of aspect are of some interest. In her earlier work (Smith 1991) she spoke of situation types that corresponded to Vendler's time schemata, and the stative situation was said to include cases such as *be tall*. In later work (Smith 2005), she divides situations into (i) Eventualities, including events and particular states, which hold at a particular time and place, (ii) General States, including generics and states that involve a pattern or regularity, and (iii) Abstract Entities, i.e., facts and propositions. In the following sentence pair, the first exemplifies an eventuality state, the second, a general state.

(24) The cat is on the mat.

(25) Peter speaks French.

But even a sentence such as (25) may hold at a particular time and place and not in another,[4] so this criterion does not seem to work well for upholding a basic distinction within the category of states. To me it appears more fruitful to postulate a general category of stative concepts that are basically atemporal but which can on the one hand be anchored to specific global states and, on the other, be the basis for the

[4] Generic and habitual sentences such as (25) have a complex temporal structure. Example (25) expresses a disposition of Peter's which is based on the activity type 'speak French'. This disposition has an extension in time which starts when Peter learns to speak French and ends when he forgets how to do it or dies. The sentence will then be true at any point in that interval.

construction of delimited states, which are temporal individuals in the same sense as events (Herweg 1991).

(26) Peter's illness began on Wednesday and ended on Sunday.

In fact, we can regard such delimited states as constructs that are defined by pairs of successive events E_1 and E_2, such that E_2 is the inverse of the E_1 and there is no other such event that follows E_1 and precedes E_2. Thus, Peter's illness, regarded as a delimited state, is defined by two events: Peter becoming ill and Peter becoming well (i.e., not ill).

5. Telic vs. Atelic

The predicates *is 22 degrees* and *rise to 22 degrees* differ in dynamicity, the former stative, the latter dynamic. Simplifying things somewhat, dynamic predicates involve change, whereas stative predicates do not. It would then be tempting to associate the time-creating capacity with dynamicity, but this would lead us astray. Instead, I want to argue that the capacity to create time constitutes a central component, or even the essence, of the notion of **telicity**, as discussed under that name, e.g., by Garey (1957), and also figuring under many other names in other works. It is perhaps most well known in the form of the distinction between activities, on the one hand, and accomplishments and achievements, on the other, in Vendler's classification (Vendler 1957).

As an example of an atelic predicate consider the Russian verb *kolebat'sja* or the English *fluctuate* as in (27–28), which could be used to describe a situation like the one in (29):

(27) Na prošloj nedele, temperatura kolebalas'.

(28) Last week the temperature fluctuated.

(29)

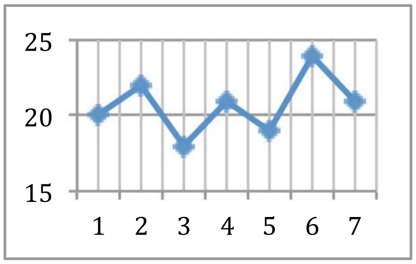

Example (29) contains a set of changes, each of which could be described by sentences involving a predicate such as *podnjat'sja* or *upast'* in Russian and *rise* or *fall* in English:

(30) Temperatura podnjalas' ot 20 do 22 gradusov.

(31) The temperature rose from 20 to 22 degrees.

(32) Temperatura upala ot 22 do 21 gradusa.

(33) The temperature fell from 22 to 21 degrees.

However, the predicates *kolebalas'* in (27) and *fluctuated* in (28) do not by themselves create time; they do not introduce a unique "before" and a unique "after". The reason is that even if we are told that there were changes in temperature, no single specific event is identified and there is no information about the number of changes or about their final result. It is often said that the telic/atelic distinction corresponds to the count/mass distinction in nouns, or alternatively that only telic predicates denote individuals in the logical sense.

But in a sense the verb *fluctuate,* like many other atelic predicates, still presupposes time in a way that makes its semantics somewhat complex. This will be clear if we look again at (28). In order to be able to verify the claim made there, we need not one observation of the temperature but a series of them; that is, we do not know whether the temperature can be said to fluctuate until we have seen it go up and down a couple of times. We could say that such a predicate has an extended verification time. But there is a further problem here which becomes clear when we consider sentences that concern on-going processes, to which we now turn.

In addition to contexts such as (28), the English verb *fluctuate* can occur in a sentence such as (34), which exemplifies a typical use of the English progressive.

(34) The temperature is fluctuating.

The English progressive is typically used to describe what is going on at a specific point in time. In the terminology of Lars Johanson (e.g., Johanson 2000) we can call this focalized time reference. Intuitively, the use of a focalized progressive usually implies that the process had been going on both before and after the point of reference. This makes the semantics of sentences such as (34) notoriously difficult. The discussion in the literature has centered around what has been called the imperfective paradox (Dowty 1979), referring to the fact that a sentence such as (35) can be true in spite of Bill's dissertation never being finished:

(35) Bill was writing his dissertation when he was hired by Google.

Example (35) contains an element of counterfactuality in that Bill probably never finished his dissertation.[5] But even quite innocent-looking progressive sentences which refer to the time of speech have a tinge if not of counterfactuality at least of speculation about the future. Compare the classical example sentence (36) to the similar-sounding (37).

(36) The temperature is ninety, but it is rising.

(37) The temperature has risen to ninety.

Examples (36) and (37) can in principle be used in the same situation, but (36) is a bolder statement in that it implies that that the temperature will continue to rise. This phenomenon does not really depend on the particular grammatical marking; the same will be true of their translations into Russian (where the verb is imperfective but not

[5] Notice that a similar kind of paradox pertains to the object of Bill's writing, his dissertation: at topic time it is not really a dissertation but probably just a draft of the first chapters, and if Bill is hired by Google it will never be anything more than that.

specifically progressive) or in Swedish (where the verb is in the simple present):

(38) Temperatura podnimaetsja.
'The temperature is rising'

(39) Temperaturen stiger.
'The temperature is rising'

Rather, it seems to be a feature that distinguishes static and dynamic sentences. For state sentences the only thing that matters is what holds at "topic time", in other words, the global state that they describe. Dynamic sentences, on the other hand, tend to go beyond that.

As with states, activity predicates can be turned into expressions denoting telic temporal individuals by being delimited. Such delimitations can take several different forms: (i) by indicating a result to be attained, e.g., *write* ⇒ *write a letter*; (ii) by a quantitative delimitation (quantizing), either by a temporal limit, e.g., *run* ⇒ *run for twenty minutes* or by some other measure, e.g., *run* ⇒ *run a kilometer*. In Vendler's terms, these are accomplishments, and in the system I am suggesting here, they are complex events, which will be further discussed below.

A special case is constituted by what can be called indefinite quantizing. This means that the activity is said to be delimited, but no precise indication of the limit is given. So-called delimitative verbs in Russian marked by the prefix *po-*, e.g., *posidet'* or *potancevat'*, are cases in point. In spite of the indefiniteness of the limit, they behave in important ways like other telic predicates, as we will see in the following section.

5.1. Narrative Progression

In the introduction, the notion of narrative progression was briefly mentioned, and I will now discuss it in somewhat greater detail. Consider the following Russian discourse segments.

(40) a. Petr otkryl dver'. Koška sidela na kovre.
'Peter opened the door. The cat sat on the mat.'

(40) b. Petr otkryl dver'. Koška vybežala.

'Peter opened the door. The cat ran out.'

It is customary to say that the sentence *Koška vybežala* in (40b), with a lexically telic verb in the perfective aspect, "moves the narration forward", whereas the corresponding stative sentence in (40a), *Koška sidela na kovre*, leaves the narration at the same point. In languages without a grammaticalized distinction between perfectivity and imperfectivity, lexical telicity would usually be sufficient to obtain the same difference between the two discourse segments (as is arguably the case for the English translations).

We may note that Russian delimitative verbs with the prefix *po-*, mentioned above, also move the narration forward, which is an argument for seeing them as telic:

(41) My posideli i ušli.

'We sat (for a while) and (then) went our way.'

The expression "moves the narration forward" seems to imply a model with a fixed timeline on which the narration moves, one step for each telic predicate. But actually the phenomenon goes beyond narratives, and it cannot always be captured by the fixed timeline concept, as we will now see.

The idea that a narration can move forward rests on the intuition that we at each point in a narration are located at a point in time.[6] This works analogously to the cursor or insertion point in a text editor, that is, the point where text is added, deleted, or changed. In a text file the default location for the insertion point is at the end of the file. Similarly, when telling a story we insert new material at the end if nothing is said to the contrary. When the user presses the return key, a new line is added to the file and the cursor moves to that line. Crucially, this is not a question of simple movement; it also involves the creation of a new element. It is thus different from pressing the down key: if you try that when the cursor is at the end of the file you won't get anywhere, since there is no line to move to. Similarly, in building a

[6] When Reichenbach (1947) introduced his notion of a point of reference, his first example was in fact taken from a narrative, although it has later been disputed whether narrative time is identifiable with this notion.

temporal representation in a narration, the default placement is at the last state description, and to get further we have to introduce a transition to a new subsequent state. In this sense narrative progression involves the creation of time as much as movement in time. The situation is different apart from a narrative context, when the discourse is anchored at the point of speech. When we hear (18), *Temperatura podnjalas' do 22 gradusov*, and accept it as true, the discourse is re-anchored to a new state of affairs, but this does not involve movement in time. Rather, what we learn is that the picture of the world that we have held so far is outdated. We could say that it is pushed back in time: the present turns into the past.

Could we mimic this in the detector model described above? When applying the rule in (14), If input > 20 and store < 20 then output "1" else do nothing, the old information is obliterated. If we want old information to be kept rather than discarded, we have to assume that the contents of storage are not simply overwritten when new information comes in. Rather, old representations are preserved like cards on a stack, with the last incoming representation on top:

(42)

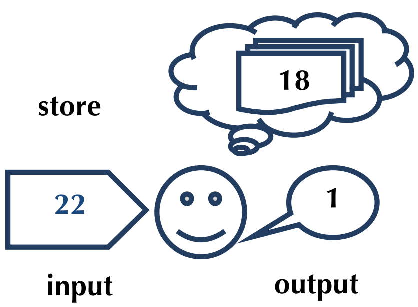

This then requires a more complex storage device than we had, e.g., in (15). Crucially, the storage in (42) has an internal sequential structure that can be interpreted in temporal terms. Interestingly, this structure is created automatically through the pushdown storage mechanism; we may say that the detector may not be time-aware, it only knows how to store information. Time-awareness would develop if we equip the device with a capability of manipulating the stored information in various ways.

5.2. The Reification Scale

We saw earlier that a representation consisting of a sequence of state descriptions and a representation consisting of a description of the initial state and a sequence of transformations, i.e., events, are equivalent. Correspondingly, the detector may store either information about states or about events or a combination of both. But storing information in terms of events rather than in terms of successive static representations means that events are reified: they are elements of the representation in their own right.

Suppose now that the detector is not only able to report on the latest incoming information, comparing it to the immediately preceding state of affairs as in (15), but also report on stored information about earlier events.

Notice that this can be seen as involving a different degree of event reification: in the first case the focus is on the comparison of two states, whereas in the second the event is treated as an entity in its own right. I would like to claim that there is in fact a scale of event reification that is reflected in the way we speak of states and events in natural language. At one end of this scale the entire focus is on the present state of affairs with no attention to the events that gave rise to it. At the other end an event is reported in isolation from states that precede or follow it. Constructions and verb forms in languages tend to occupy a larger or smaller piece of the scale. For instance in English:

(43)

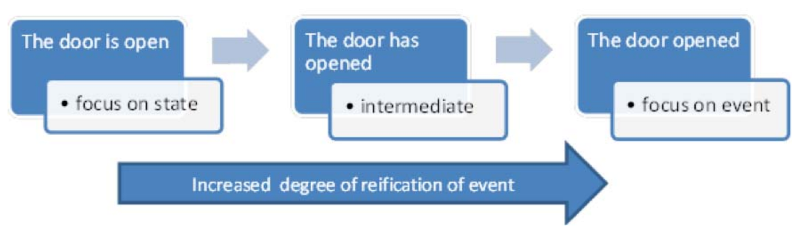

In other words reification would not be an all-or-none phenomenon, but rather a graded one, depending on the attention we pay to an entity. Differences in the degree of reification may help account for distinctions between grammatical entities such as the past tense and the perfect in English.

6. Observation

Taken by itself, a state description is atemporal: by definition it does not contain any temporal relationships between its elements. However, a state description may be anchored in time, and often it has to be in order to make sense. One important way in which an anchoring is brought about is through an observation. For instance, I can make an observation of the temperature in a certain location by reading it off a thermometer. The reading of the thermometer is in itself an event with a unique before and a unique after. This adds an element of apparent paradox into the system, since the state will normally have a certain duration, although the reading, or rather the completion of it, is seen as momentaneous. From the linguistic point of view, this means among other things that descriptions of states are often combined with temporal expressions identifying a point in time, which is the point in time when the state is observed (in German-language literature the term *Betrachtzeit* is sometimes used for this). Another effect is that a description of a state can be integrated into a sequence of events, e.g., in a narrative text, as in (44):

(44) a. John opened the door.
b. The room was pitch dark.
c. He turned on the light.

Sequences like (44) have been subject to considerable discussion in the literature. It is sometimes claimed that a state such as the one reported in (44b) necessarily holds when the event reported in (44a) takes place.

However, Hinrichs (1986), and following him Boogaart (1999: 104) assume that a telic predicate such as the one in (44a) introduces a temporal reference point or point of perspective just after the event of opening the door. Boogaart notes in this connection that "it is not part of the semantic content" that the state in (44b) overlaps temporally with the event in (44a). This can be seen from examples such as (45) (given that the lamp in a fridge normally turns on automatically when the door is opened):

(45) John opened the door of the fridge. The inside was brightly lit.

But we can take this one step further. Given that the normal interpretation of sequences such as (44) is that the state referred to by (44b) is observed by the protagonist of the story, we can see (44) as representing a sequence of three events, where the second one is the event of Peter perceiving or becoming aware of the room being pitch dark—something which is itself a state description. The result could be depicted as:

(46)

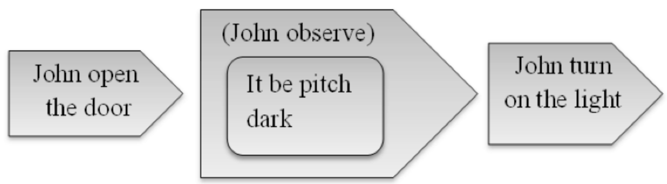

Theories of aspect[7] often rest upon a notion of viewpoint, which the aspect chosen when speaking of a temporal entity is said to depend on. As we can see in (46), such a viewpoint can actually be part of the narrative event structure itself.

Hinrichs (1986: 67), while acknowledging that considerations of ontological purity might incline one to treat states and events as fundamentally different entities, argues that "this would unduly complicate the system", as it would force us to simultaneously build two distinct structures for a discourse. To this may be objected that descriptions in terms of successive states and descriptions in terms of events are not distinct since they are largely interdefinable. But as the preceding example shows, treating them as different entities does not necessarily prevent us from integrating them into one and the same structure.

6.1. Simple and Complex Events

A transition, as defined above, represents the most elementary type of event. A transition need not be instantaneous in physical time, but relative to the model it can still be regarded as simple in the sense that it

[7] Obviously, not only theories of aspect. That a story can be told from a participant's viewpoint is basic to any theory of narration. My point here is the interaction of viewpoint with narrative progression.

relates to adjacent states. Suppose we have four observations of the temperature at a certain location, as in the graph below.

(47)

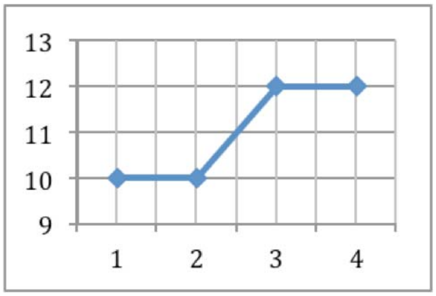

Between observations 2 and 3, the temperature rose by two degrees, from 10 to 12. From the point of view of (47), this is a simple transition. But consider what happens if we have an intermediate observation at which the temperature is registered as 11 degrees.

(48)

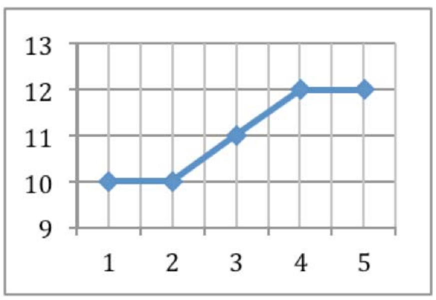

We can now describe what happens in terms of two simple events, each consisting of an increase by one degree. However, it is also possible to continue to speak of one complex event that increases the temperature by two degrees, but which involves an intermediate stage at which the temperature has risen by one degree only.

A distinction is often made between punctual and non-punctual events, and correspondingly between punctual and non-punctual verbs (verb phrases, predications). In the Vendlerian tradition, this is the distinction between **achievements** and **accomplishments**. It seems to me that what is at stake is rather the distinction between simple and complex events, that is, whether an event is thought of as having intermediate stages or not. This is often a matter of choice, of the granularity of the description[8].

In the examples above, the complex event or accomplishment could be seen as a series of simple events, involving stepwise changes of a single parameter. But the typical accomplishments that we talk of in everyday life have a less straightforward structure. In particular, this is true when the intentions of an agent are involved. Consider for instance the action of baking a cake. It consists of getting the ingredients out of the larder, mixing them into a batter, pouring the batter into a baking pan, turning the oven on, putting the pan in, and getting it out after the appropriate time. What keeps all these events together is that they all have the same goal: to create the result, that is, the cake. Only actions that are done with this goal are unequivocally part of the complex event of baking a cake.[9] The end state is thus crucial for the understanding of a telic predicate, whether complex or not. It is really the raison d'être or point of the whole thing. There are in fact two possible perspectives on accomplishments: they could be seen either as an activity followed by a boundary, where the activity is basic, or as an event preceded by whatever led up to that event, in which case the event is basic. In this article I have generally tried to choose the latter option, whereas many approaches to temporal semantics tend to choose the former. I think that is a natural consequence of trying to fit all eventualities into the same mold or template, where they are all seen as being built up by the same components or phases, although for some types some of the phases are missing. Accomplishments will then be seen as the most complete eventuality type, with the others all

[8] In cognitive linguistics, granularity figures under the name of ”specificity” as one of the dimensions of ”construal”, and it has been suggested to me that it would be appropriate to use here the latter term, which is defined (Langacker 2008: 43) as ”our manifest ability to conceive and portray the same situation in alternate ways”. I am somewhat hesitant to do so, however, since I find the distinction between construal and what Langacker calls content (here figuring as the same situation) rather problematic. In fact, Langacker himself mentions specificity as an example of the distinction ”not at all” being a sharp one.

[9] Mehlig (2012) points out that it is a characteristic of telic predicates denoting what I here call complex events that the sub-events have to occur in a particular order. For instance, you have to make the batter before putting it in the oven. But if it is a condition on the sub-events included in an accomplishment that they are performed in order to reach a certain goal, this follows automatically.

being deficient in one way or other. However we look at it, an accomplishment will be concluded by a simple transition event, i.e., an achievement. Moreover, this event should be uniquely identifiable, since without that condition we would be dealing with an activity.

Consider now the predicates *bake a cake* and *finish a cake*. These denote events that have the same result state, the existence of a cake, but they differ in what they include. As we just saw, baking a cake includes all the actions that have the creation of the cake as goal, while finishing a cake just involves the final point of the process. The first then is a complex event or an accomplishment, whereas the second is a simple event or achievement. However, the distinction is far from cut and dried: a verb like *finish* may also be used of the final stages of a process, rather than just the final point. Note also that *finish* normally presupposes the existence of a longer process, which would be treated as an accomplishment in a Vendlerian system. Actually, it may be slightly unfortunate that Vendler used *reach the top* as his primary example of an achievement, as it belongs to the rather special class of predicates that are connected with such a presupposition (before you reach the top you will normally have been climbing for quite a while). This detracts from the central role that simple transition events must have in any temporal ontology, as was recognized by von Wright in his logic of change.

7. Time Measurement

This brings us to the important question of time measurement. As humans, we are able to make intuitive judgments about the length of periods of time. We are also able to measure time in a more objective way by using natural phenomena such as the movements of celestial bodies and technological devices such as clocks. In the latter case, we normally assess the length of a period by relating it to a series of events, usually of the same type. We know that a period lasts for five days if it begins at the first and ends at the sixth in a series of sunrise events. A mechanical or digital clock is constructed to move forward at a certain speed or with certain intervals. Notice that measuring time by counting recurring events does not necessarily mean that these events occur at even intervals from each other, as is shown by expressions such as *two beers later*. But not even time in calendarical units such as months and years are always of the same length. And yet the existence of leap years does not keep us from measuring time in years. Furthermore, there has to be some ultimate series of events that are assumed to take place at equal intervals. This has to remain an assumption because otherwise we would have an infinite regress. Thus, a second is defined in the international SI system as "the duration of 9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the caesium 133 atom".

As was mentioned, humans are able to assess the length of time intervals even without recourse to an external clock. This presupposes, however, that we have some kind of internal clock that the assessment relies on. The mechanisms behind this are not well understood, although we know that all living creatures have some kind of biological clocks that rely both on external stimuli and internal mechanisms, most importantly on those that regulate the daily cycle. Friedman (1993) stresses the importance of the daily cycle for positioning events in memory. Interestingly, it is often easier to remember at what time of day an event took place than to remember what day it happened on.[10]

One point to be made here is that assessment of the length of intervals appears at least on the surface to be continuous rather than discrete, if it is not supported by knowledge about external events. If I remember that I was awake in the night, I can feel that I was awake for a long or a short while, but I will not be able to say that I was awake for exactly ten minutes. On the other hand, such assessments tend to be made of intervals that are themselves delimited by events, such as, in this case, waking up and falling asleep again. Most periods or intervals in time that we speak of or even have names for are in fact definable in a similar way. A day can be seen as a delimited state characterized by a sunrise and and a sunset event (or in many cultures, two sunset events). But any two events that are not simultaneous could be used to define an interval.

Intervals, in such a system, are a derived notion based on events. Likewise, the notion of an event taking place within an interval can be defined by the relations between that event and the pair of events that define the interval. That an event took place on a certain day means that it succeeds the beginning of that day and precedes its end.

---

[10] See also Fauconnier and Turner 2002 on the role of the daily cycle in the network from which time emerges.

8. Conclusion

In this paper, I have tried to find an alternative to the common approach to temporal semantics, where every sentence or predicate in discourse is supposed to correspond to an eventuality to be placed on a timeline. With inspiration in earlier work of Georg Henrik von Wright, Michael Herweg, and Sebastian Löbner, I wanted to see if an approach that takes change as basic for time could be used to build cognitively realistic temporal representations.

The basic elements of the model are global states, which are timeless taken in isolation but connected by transition events, which transform one global state into another and thereby define the temporal relationships between them. Transition events, corresponding to Vendlerian achievements, represent simple changes which are then the basis for all other constructs in the model, most notably delimited states, Vendlerian activities (atelic dynamic eventualities), and accomplishments (telic non-punctual eventualities), but also time points and intervals. Transition events are further instrumental in constructing narrative structures and are responsible for narrative progression.

References

Boogaart, Roony J. U. (1999) *Aspect and temporal ordering: A contrastive analysis of Dutch and English.* The Hague: Holland Academic Graphics. [*LOT International Series*, 27.]

Cassandras, Christos G. and Stéphane Lafortune. (2008) *Introduction to discrete event systems.* New York: Springer.

Dahl, Östen. (1981) "On the definition of the telic-atelic (bounded-nonbounded) distinction". Philip J. Tedeschi and Annie E. Zaenen, eds. *Tense and aspect.* New York: Academic Press, 79–90.

———. (1985) *Tense and aspect systems*. Oxford: Blackwell.

———. (2007) "Towards an ecological semantics of tense and aspect". Daniele Monticelli and Anu Treikelder, eds. *L'aspect dans les langues les théories: Similitudes et différences*. Tartu: Tartu Ülikool, 111–24.

Davidson, Donald L. (1967) "The logical form of action sentences". Nicholas Rescher, ed. *The logic of decision and action*. Pittsburgh: University of Pittsburgh Press, 81–120.

Dowty, David R. (1979) *Word meaning and Montague grammar: the semantics of verbs and times in generative semantics and in Montague's PTQ*. Dordrecht: Reidel.

Fauconnier, Gilles and Mark Turner. (2008) "Rethinking metaphor". Raymond W. Gibbs, Jr. ed. *Cambridge handbook of metaphor and thought*. Cambridge: Cambridge University Press, 53–66.

Fleischman, Suzanne. (1990) *Tense and narrativity: From medieval performance to modern fiction*. London: Routledge.

Friedman, William J. (1993) "Memory for the time of past events". *Psychological bulletin* 113(1): 44.

Garey, Howard B. (1957) "Verbal aspect in French". *Language* 33(2): 91–110.

Herweg, Michael. (1991) "A critical examination of two classical approaches to aspect". *Journal of semantics* 8(4): 363–402.

Hinrichs, Erhard. (1986) "Temporal anaphora in discourses of English". *Linguistics and philosophy* 9(1): 63–82.

Holland, John H. (1995) *Hidden order: How adaptation builds complexity*. Reading, MA: Addison-Wesley.

Johanson, Lars. (2000) "Viewpoint operators in European languages". Östen Dahl, ed. *Tense and aspect in the languages of Europe*. Berlin: Mouton de Gruyter, 27–188.

Kamp, Hans, Josef van Genabith, and Uwe Reyle. (2011) "Discourse Representation Theory". Dov M. Gabbay and Franz Guenthner, eds. *Handbook of philosophical logic*. Vol. 15. Amsterdam: Springer, 125–394.

Langacker, Ronald W. (2008) *Cognitive grammar: A basic introduction*. Oxford: Oxford University Press.

Löbner, Sebastian. (1988) "Ansätze zu einer integralen semantischen Theorie von Tempus, Aspekt und Aktionsarten". Veronika Ehrich and Heinz Vater, eds., *Temporalsemantik: Beiträge zur Linguistik der Zeitreferenz*. Berlin: Walter de Gruyter, 163–91.

Quine, Willard van Orman. (1985) "Events and reification". Ernest LePore and Brian P. McLaughlin, eds., *Actions and events: Perspectives on the philosophy of Donald Davidson*. Oxford: Blackwell, 162–71.

Reichenbach, Hans. (1947) *Elements of symbolic logic*. New York: Macmillan.

Rosenberg, Daniel and Anthony Grafton. (2010) *Cartographies of time: A history of the timeline*. New York: Princeton Architectural Press.

Smith, Carlota S. (1991) *The parameter of aspect*. Dordrecht: Kluwer Academic Publishers.

———. (2005) "Aspectual entities and tense in discourse". Paula M. Kempchinsky and Roumyana Slabakova, eds. *Aspectual inquiries*. Dordrecht: Springer, 223–37.

Vendler, Zeno. (1957) "Verbs and times". *The philosophical review* 66(2): 143–60.

von Wright, Georg Henrik. (1963) *Norm and action: A logical inquiry*. London: Routledge and Kegen Paul PLC.

Department of Linguistics
University of Stockholm
106 91 Stockholm
Sweden
oesten@ling.su.se

Received: August 2012
Revised: January 2013