

Iconic Update

Katja Jasinskaja

Abstract

This paper presents an idea of how the notion of iconicity, in particular, temporal iconicity in narrative discourse, can be made explicit in a formal semantic framework, such as the DRT. It argues for a distinction between two discourse update modes: (a) non-iconic update, which characterises the standard way utterances are interpreted; and (b) *iconic update*, which mimics the process of direct perception. Iconic update plays a central role in the account of reportive and historical present, as well as the temporal progression in past tense narrative. The proposed approach reconstructs the basic predictions of previous accounts of these phenomena, and allows for extensions to new empirical domains that go beyond temporality.

Streszczenie

W artykule pokazuje się, jak pojęcie ikonizmu, w szczególności ikonizmu temporalnego może zostać eksplikowane w takiej formalno-semantycznej teorii, jaką jest Teoria Reprezentacji Dyskursu. Przedstawiono argumenty na rzecz rozróżnienia pomiędzy dwoma sposobami uaktualnienia w dyskursie: (a) nie ikonycznym uaktualnieniem, które jest standardową drogą interpretacji wypowiedzi, a (b) *ikonicznym uaktualnieniem*, które naśladuje proces bezpośredniej obserwacji. Ikoniczne uaktualnienie odgrywa centralną rolę w wyjaśnieniu „reportażowego” oraz „historycznego” czasu teraźniejszego, jak również progresji czasowej w narracji w czasie przeszłym. Proponowane nie posiada taką samą moc eksplanatoryczną jak wcześniejsze podejścia, a zarazem może być zastosowane w innych, m.in. pozaczasowych obrębach empirycznych.

1 Temporal iconicity in discourse

The focus of this paper is on the well-known tendency of the textual order of utterances in narrative discourse to mirror the temporal order of events described. It is a widely spread assumption that this phenomenon is one of the numerous manifestations of *iconicity* in language, i.e. cases where the relationship between the form and the meaning of a linguistic sign is that of resemblance. C. S. Peirce's characterisation of signs according to the form-meaning relationship was introduced into linguistics by Roman Jakobson, along with the temporal iconicity idea:¹

¹See Nöth (2001) for an overview on the notion of 'iconicity' in semiotics, linguistics, and literature.

The chain of verbs—*Veni, vidi, vici*—informs us about the order of Caesar’s deeds first and foremost because the sequence of co-ordinate preterits is used to reproduce the succession of reported occurrences. The temporal order of speech events tends to mirror the order of narrated events in time... (Jakobson, 1971)

It is undisputed that the narrative presentation order is, in one way or another, given by the nature itself. That is, some inherently linguistic characteristics of this discourse type reflect something inherently non-linguistic. However, views differ on the question of level at which this similarity should be placed. In the classical view inspired by Peirce and Jakobson, the relationship is defined as resemblance between form and meaning. A more psychological take on the matter is to say that the process of language comprehension can be conceived of as “vicarious experiencing of events in the real world” (Segal, 1995; Zwaan, 1999; Zwaan et al., 2001). In other words, resemblance is established between the *processes* of language comprehension and perception of non-linguistic input. When speakers narrate events in the chronological order they exploit the hearers’ skill to interpret events in the real world. The observed resemblance between form and meaning with respect to the order of utterances/events is a consequence of similarity in the processes.

Even though the basic intuition behind temporal iconicity is also shared by formal semanticists, the way chronological order is standardly implemented in their theories does not reflect the idea of resemblance in any way, be that at the level of meaning-form relationship or at the level of the process. Certain tense/aspect features (or feature combinations) are treated as anaphors that refer to events or times introduced previously in the discourse and establish a ‘later than’ relation between their referent and the time of the main event described in the current sentence. Or put differently, sentences come with an implicit “after that” in their semantics (Kamp and Reyle, 1993; Kamp et al., 2005).

The goal of this paper is to sketch out a possible avenue for a formal conception of temporal iconicity as resemblance, within a dynamic semantic framework, such as the Discourse Representation Theory (DRT, Kamp and Reyle, 1993; Kamp et al., 2005). The meaning of a sentence in dynamic semantics is its context change potential, i.e. a function that maps the shared information state of the communication participants before the sentence was processed to the information state after. The change of the information state as each sentence is being processed is also referred to as discourse *update*. The change in beliefs of an agent that results from direct perception of changes in the world can also be conceived of as a kind of update, using the same formal tools. In this paper I will argue that update from linguistic communication and update from direct perception differ in the way they treat the time of update, the *now*. In discourse update, *now* is (normally) an extended time interval comprising the whole discourse, and in that sense, constant within the boundaries of the same discourse. In update from direct perception, there is a new *now* for each perceived change in the world. I further define the notion of *iconic update*, an update mode used in communication which mimes update from

direct perception in the way it treats the update time. In other words, iconicity is placed at the level of the process—the process of update. In addition to an account of temporal iconicity in normal, past tense narrative, this approach allows for a simpler theory of tense, and provides a framework that naturally accommodates a host of “narrative” phenomena: reportive and historical present, as well as other manifestations of perspective shifting.

After a short introduction to the DRT account of chronological interpretation of the narrative in section 2, section 3 presents an extension of the standard notion of update to perception, as well as the notion of iconic update and the account of chronological interpretation based on it. Section 4 explores some possible applications of the proposed approach that go beyond temporality. Section 5 concludes with directions for further research.

2 Narration in DRT

The account of tense and aspect in narrative discourse in the framework of DRT² concentrates on facts such as the contrast between (1) and (2). The basic observation is that events push the narrative forward, whereas states do not. For instance, both sentences in (1) describe events (achievements or accomplishments), i.e. either instantaneous changes in the state of the world, or changes that culminate and reach completion at a particular instant. Thus in (1), we interpret the second event as happening later than the first. In (2), the second sentence describes the same act of pulling a gun from its holster, but viewed “from within”, as a process stretching over a period of time which does not reach completion within that period. For the present purposes, processes belong to the same category as proper states (e.g. *The door was closed*), i.e. the category of states in a broad sense, or durative eventualities. The generalisation is that states typically overlap the last mentioned event, i.e. in (2) Josef turns around while the man is pulling his gun.³

- (1) a. Josef turned around.
b. The man pulled his gun from its holster.
- (2) a. Josef turned around.
b. The man was pulling his gun from its holster.

The DRT account of (1) goes roughly as follows: We assume that the communication participants start with an empty common ground, i.e. an empty set of mutual beliefs (this is absolutely unrealistic, but harmless for the task at hand). After processing the first sentence, the common ground is updated with the new information and reaches the state captured (with numerous simplifications) in the discourse

²Here I am recapitulating the introduction in Kamp et al. (2005), but the original ideas behind this approach go back to Partee (1984) and Hinrichs (1986).

³In this particular example the difference in interpretation is effected by the choice of aspectual form, simple past in (1) vs. past progressive in (2). However, the state/event distinction is influenced by many other factors, for instance, the intrinsic aspectual class of the verb (see e.g. Dowty, 1986).

representation structure (DRS) in (3): e_1 is a variable that refers to the event of Josef turning around; t_1 is the *location time* of that event—the event is included, i.e. occurs within t_1 ($e_1 \subseteq t_1$); the contribution of the past tense is the condition $t_1 \prec now$, which says that the location time of the event strictly precedes the speech time *now*.

$$(3) \quad \boxed{\begin{array}{l} e_1, t_1 \\ e_1 : \text{'Josef turn around'} \\ e_1 \subseteq t_1 \\ t_1 \prec now \end{array}}$$

The semantics of the sentence (1b) is represented in (4). It is an ordered pair that consists of a set of presuppositions (the set of DRSs in the curly brackets), and the main DRS representing the assertive content of the sentence. The main DRS has the same structure as that in (3), introducing a location time t_2 before *now* and an event e_2 (of a man pulling his gun) included in its location time.

$$(4) \quad \left\langle \left\{ \boxed{\begin{array}{l} r \\ r \prec t_2 \end{array}} \right\}, \boxed{\begin{array}{l} e_2, t_2 \\ e_2 : \text{'Man pull gun'} \\ e_2 \subseteq t_2 \\ t_2 \prec now \end{array}} \right\rangle$$

The anaphoric nature of tense is captured by the presupposition on the left hand side of (4), which imposes a condition on the common ground reached so far that it should contain a (temporal) referent r (reference time) lying earlier than the location time t_2 of the new event. We do find an appropriate referent in the context DRS (3)— t_1 , the location time of the first event—but its relation to t_2 is not specified. Since nothing contradicts the assumption that $t_1 \prec t_2$, this part of the presupposition gets accommodated. The resulting DRS is then updated with the content of the main DRS in (4), yielding the representation in (5).

$$(5) \quad \boxed{\begin{array}{l} e_1, e_2, t_1, t_2 \\ e_1 : \text{'Josef turn around'} \\ e_1 \subseteq t_1 \\ e_2 : \text{'Man pull gun'} \\ e_2 \subseteq t_2 \\ t_1 \prec t_2 \prec now \end{array}}$$

The presupposition in (4) encodes basically the same condition as the expression *after that*. It is as if event sentences always came with an implicit *after that* in their semantics.

For comparison, the semantics of the sentence (2b), which due to the use of past progressive presents the eventuality as an ongoing process, differs from the semantics of (1b) in two respects. First, while events are included in their location times ($e_1 \subseteq t_1$), the reverse relation holds for states, which include their location

times: $t_2 \subseteq s_2$. This reflects the idea that states and processes are viewed “from within”. Second, states come with a different presupposition: the location time t_2 of the state is equal to a contextually given reference time r , which is, again, resolved to the location time t_1 of the previously mentioned event e_1 . This implements the idea that states do not “push the narrative forward”. It follows that the state s_2 overlaps the event e_1 : $e_1 \subseteq s_2$.

This simplistic version of the approach has a number of well-known problems. First of all, if the *later than* relation is encoded directly in the semantics of simple past (or arrived at by semantic composition of simple past with the intrinsic aspectual class of the verb, etc.) then it should be present in all its uses. However, the chronological interpretation is restricted to narration, and there is a whole range of non-narrative uses of the past tense which do not show temporal iconicity. In (6), the temporal order of the two events is not specified—they could have occurred in any order or simultaneously. It seems that the temporal order is simply not the point in this discourse. In (7) the second event is interpreted as the cause of the first, therefore the reverse temporal order is inferred.

(6) Max spilt a bucket of water. John dropped a jar of cookies.

(7) Max fell. John pushed him.

Even if we assumed that event sentences in simple past are ambiguous between a version with the *after that* presupposition (in the narrative) and one without (in other contexts), it looks like a mere coincidence that the relation is *after that*, and not, let’s say *before that*, which would give rise to narratives with reverse chronological order. Obviously, this can be explained by assuming that grammar is optimised to express the natural order of events. However, the relation between the natural order of events and the order of presentation in the narrative is not made explicit in the theory in any way.

Concerning the problem in (6) and (7), nowadays it has become more or less standard to assume that tenses like the English simple past do not encode temporal relations to previously mentioned events. According to segmented DRT (SDRT, Asher and Lascarides, 2003), aspect, world knowledge, discourse connectives and other characteristics of the utterances and the speech situation jointly operate as premises in defeasible inference of *rhetorical relations*, such as *Narration* in (1), *Background* in (2), *Parallel* in (6) and *Explanation* in (7), which, in turn, determine (or not) the temporal relation: *Narration* establishes succession, *Explanation* establishes the reverse temporal order, and *Parallel* leaves the relation unspecified. However, existing SDRT-based proposals, just like their DRT predecessors, do not go as far as modelling temporal sequentiality as a consequence of the fact that events in a narrative are presented *as if* happening before the eyes of an observer. This paper makes first steps towards such a model, by characterising the way in which observers perceive events, and stating explicitly, in which way the processes of direct observation and discourse interpretation are similar.

3 Update time

To begin with, two remarks about the nature of discourse representation structures in DRT are in order. First, they are thought of as *conceptual belief representations* that result from the interpretation of utterances, on the one hand, and have themselves a model-theoretic interpretation, on the other. That is, they mediate between language and the real world, but they are not themselves part of the real world, but rather of the speakers' representation of the world. Second, as is standard in dynamic semantics, the target DRS that is updated and accumulates information from incoming utterances, such as (5) above, represents the *common ground* of communication participants, i.e. their mutual beliefs shared in the process of communication. However, nothing in the definition of DRSs prevents them from being used to represent belief states in general, no matter whether those beliefs are acquired by communication or from direct observation of events in the world.⁴ Similarly, the notion of update can be generalised to include both update by communication and update from direct perception. Update time, in turn, is the generalised notion for speech, or utterance time, on the one hand, and direct perception time, on the other.⁵ The central idea of the present approach is that update time is treated differently during normal (non-iconic) update by communication, and during update by direct perception, which in turn is imitated in iconic update. This section takes up the issue of time in these three update modes.

3.1 Update by communication

The primary function of tense is to relate the reported eventuality to the utterance time *now*—this much is clear. But what is the utterance time? For instance, do the utterances in (8) have the same or two distinct utterance times?

- (8) a. Josef turned around.
b. The man pulled his gun from its holster.

⁴In that case, it would be more natural to refer to them as *belief representation structures*, rather than *discourse representation structures*, since there need not be any discourse going on. Nevertheless, I will stick to the received acronym 'DRS' referring to these structures.

⁵Strictly speaking, one should distinguish between four times: utterance time, utterance perception time, event time, and event perception time. It would be natural to reserve the term 'update time' for utterance perception and event perception time, whereas the generalised notion for utterance and event time could be something like 'update-triggering event time', i.e. the time of event that causes the update of an agent's belief state—an utterance in the case of communication, and the observed event in the case of direct perception. The difference between the event time and the event perception time is normally negligible: it takes time for light to reach retina and for the two-dimensional activation pattern on the retina to undergo several processing stages before it is mapped to a concept like 'spill', 'water', etc. (cf. Bennett et al., 1989, p. 4: "Perception is a process of inference."). However, these processes are fast and automatic, so the update time can be safely assumed to tightly follow the event time. With utterances this is more of an issue, especially in writing, since as we all know, a text can be written centuries before it is read. In this paper I will follow the widely spread tradition of ignoring this difference.

The standard assumption, also reflected in the DRS representation (5) in the previous section, is that the utterances have the same *now*, which corresponds to an extended time interval which normally comprises all the utterances of the same discourse (see e.g. discussion in Reyle et al., 2007, pp. 607–609). This is what allows us to infer, for instance, that the event of the man pulling his gun from its holster occurred not only before the utterance (8b), but also before the utterance (8a).

If this were not so and each utterance had its own *now*, then in order to answer a simple question like *What’s the weather like?* one would have to use the past tense: *It was raining*, because the speaker asking the question would be asking about the time of *his* utterance, which lies in the past by the time the answer is given. This is obviously not the way we normally use language. When we talk about the present, we are normally interested in facts that will hold for some reasonable period of time, at least as long as we are engaged in the conversation, so one can build up common ground and take those facts for granted once they are established between the speakers.

Another manifestation of this tendency in natural language, is the fact that perfective aspect does not combine easily with present tense. In English, for instance, non-stative predicates are awkward in simple present, if this tense is used to express simultaneity with the time of utterance, cf. Zucchi (2005) and example (9) adapted from his example.⁶ Similarly, present tense morphology on Russian perfective verbs only has interpretations referring to the future.

- (9) a. # At this moment, John attends the Tense Colloquium in Paris.
 b. At this moment, John is attending the Tense Colloquium in Paris.

The generalisation is well-known (see e.g. also Smith, 1991; Kamp and Reyle, 1993). Reyle et al. (2007, pp. 607–609) formulate it as a prohibition of the $ev \subseteq now$ relation for the normal (non-reportive) uses of the present. Since non-stative predicates with the English simple tenses, as well as Russian perfective verbs describe events, and events are included in their location times, $e \subseteq t$, those aspectual forms are ruled out when $t = now$. Reyle et al. (2007, p. 608) conjecture that this is due to the assumption that: “nothing that is of importance to what it [the discourse] talks about changes while it is in progress”. Time stops when we talk. Reyle et al. suggest that this is ultimately due to the fact that *now* is treated as a point, and that all instances in the duration T_d of the discourse are thought of as equally good candidates for playing the role of utterance time. The assumption that *now* is a point does not seem necessary (and is even potentially problematic) for the present analysis which is supposed to span across both the normal uses of the present tense (under normal non-iconic update) and the “marked” (reportive, historical) uses (under iconic update). However, it seems that Reyle et al.’s gener-

⁶That is, unless we have to do with reportive or historical present, which belong to the group of cases, where we are not talking about the present in the normal way and the generalisations presented in this section do not apply. Reportive and historical present will be our primary motivation for the notion of iconic update in section 3.3.

alisation can also be explained by the pragmatics of communication. Under normal circumstances, speakers follow the silent agreement not to talk about things that are changing now, as they are talking, because in that case, whatever has been said and agreed upon between the speakers cannot be taken for granted anymore at a later stage in the conversation. In other words, the common ground gets unstable and unreliable. This is not to say that people cannot talk about things while they are happening. It is just that that requires a special form of update of the common ground, the iconic update, discussed in section 3.3.

Finally, what does this view of update and this conception of update time predict with respect to the temporal interpretation of connected discourses? If we remove the presupposition of a relation to a contextually given reference time from the semantics of the sentences (for reasons given in the previous section), then we get no predictions whatsoever. In (10) and (11), for instance, tense semantics locates both events before *now*, i.e. before the beginning of the discourse, but there is no constraint on the order of the events with respect to each other. This is exactly what we want for (10), and it is consistent with the reverse temporal interpretation of (11). (The actual inference of the reverse temporal order in (11) is driven by independent mechanisms, see e.g. Jasinskaja and Karagjosova (2011), and will not be discussed in this paper.)

- (10) a. Max spilt a bucket of water.
 b. John dropped a jar of cookies.

- (11) a. Max fell.
 b. John pushed him.

3.2 Update from direct observation

The way we treat the update time in direct perception is quite different. When we perceive events as they are happening, the relationship between the update time and the event time is much more direct. Events are not marked for tense. Rather, every event happens *now*, and we cannot directly observe past or future events. That is, the moment when the information of the event enters our mind (the update time) is the only handle we get on the actual event time. In this case it makes more sense to treat all update times as distinct. Otherwise, we would not be able to store the observed events in the right temporal order.

Suppose the observer directly perceives the sequence of events described in (12). These are not utterances s/he hears, but metalanguage descriptions of the events s/he observes. Figures 1 and 2 show conceptual representations of the observed events as they enter the overall conceptual representation of the observer's current belief state. One can think of them as the result of interpretation and conceptualisation of the visual, auditory, and other sensory input effected by the actual happenings. The condition $e_n \subseteq \textit{now}$, which could perhaps even be strengthened to $e_n = \textit{now}$, reflects the idea that the event time *is* the update time.

- (12) a. Max spills a bucket of water. [on Dec 16, 2011 AD, at 16:04:53]
 b. John drops a jar of cookies. [on Dec 16, 2011 AD, at 16:05:09]

Let us ignore all the previous beliefs of the observer and start with an empty belief state representation. The update of an empty belief state with the DRS in Figure 1, results in a DRS equivalent to that in Figure 1, which then represents the “context” to be updated with the DRS in Figure 2 by the time the second event is processed. If the observer were using the standard communicative update procedure described in previous sections, the second update would either result in the inference that the events happen simultaneously (on the assumption that $e_1 = now$ and $e_2 = now$) or in no temporal inference at all (on the assumption that $e_1 \subseteq now$ and $e_2 \subseteq now$), that is, after perceiving the two events happening strictly in a sequence, the observer would still not know in which order they happened.

One way to distinguish between the two instances of *now*, is to put the first instance in the scope of an appropriate operator. For this purpose, we introduce the operator MEM (memory), which is supposed to reflect the idea that with each new observation, the result of the previous update of the observer’s belief state becomes *a memory of an earlier perception*. The result is shown in Figure 3.

The definition of the MEM operator is given in (13), using Kamp et al.’s (2005) framework for propositional attitudes. The DRS K representing the previous belief state is embedded in an attitude context—a perception state s^{att} whose agent is the observer him/herself (i , for ‘I’, 1SG). The attitude state is located in the past, i.e. before the new, post-observation now_2 : $t \subseteq s^{att}$; $t \prec now_2$. All occurrences of now_1 in K are internal to the attitude context, and are mapped to its external anchor t . That is, now_1 represents the observer’s ‘now’ within that past memory, whereas now_2 represents the actual, new ‘now’.

$$(13) \text{ MEM}(K) = \boxed{\begin{array}{l} t, s^{att} \\ s^{att} : \text{ATT}(i, \langle \text{PERCEIVE}(K) \rangle, \langle now_1, t \rangle) \\ t \subseteq s^{att} \\ t \prec now_2 \end{array}}$$

If we unpack the content of the attitude state, mapping all the internal anchors to their external counterparts, the old, pre-observation ‘now’ becomes past, i.e. preceding the new, post-observation ‘now’, as shown in Figure 4.

In sum, we perceive changes in the environment one after the other, these enter our conceptual representation as they are perceived, so the event time out there in the real world gets mapped more or less directly to a time in the conceptual representation. One might wonder, what we really know about perception of events and how realistic this strictly linear view of perception is. Obviously, this view is a simplification. Even if we restrict our attention to vision alone: Continuous observation might be conceived of as a linear sequence of fixations—states in which the observer focuses on a particular fragment of his/her visual field, and in which the lion’s share of visual information is consumed. However, the perception of one even relatively simple event will often take several fixations, which are not

| |
|----------------------------------|
| e_1 |
| $e_1 : \text{'Max spill water'}$ |
| $e_1 \subseteq \textit{now}$ |

Figure 1: DRS for (12a).

| |
|------------------------------------|
| e_2 |
| $e_2 : \text{'John drop cookies'}$ |
| $e_2 \subseteq \textit{now}$ |

Figure 2: DRS for (12b).

| |
|------------------------------------|
| e_2 |
| MEM |
| e_1 |
| $e_1 : \text{'Max spill water'}$ |
| $e_1 \subseteq \textit{now}_1$ |
| $e_2 : \text{'John drop cookies'}$ |
| $e_2 \subseteq \textit{now}_2$ |

Figure 3: Update from direct observation

| |
|------------------------------------|
| e_1, e_2, t_1 |
| $e_1 : \text{'Max spill water'}$ |
| $e_1 \subseteq t_1$ |
| $t_1 \prec \textit{now}$ |
| $e_2 : \text{'John drop cookies'}$ |
| $e_2 \subseteq \textit{now}$ |

Figure 4: The content of MEM unpacked.

guaranteed to be contiguous and ordered with respect to the perception of other events. This is particularly true for dynamic concepts (e.g. motion), which require observation over a stretch of time, cf. Papafragou et al. (2008), Schmiedtová and Sahonenko (2008), Schmiedtová et al. (to appear). That is, perception of an event has its own duration, and the order of perception is not necessarily strictly linear. We will ignore these complications for the time being, assuming that at a certain relatively coarse level of granularity the linear view of perception is good enough: what happens first, is perceived first and reaches first the conceptual representation.

3.3 Iconic update

Iconic update is an update mode used in communication, which, however, imitates update from direct observation in the way it treats the update time. A prerequisite for the use of iconic update is the presence of an *observer*, which can be the speaker himself, another real individual, or a fictitious observer. In our DRS language, this condition will be satisfied by the presence of an operator OBS defined in (14), which means that there must be an individual o that perceives some state of affairs described in the DRS K at time t^o (observation time). All references to the update time *now* in K are mapped to their external anchor, the observation time t^o .

$$(14) \text{ OBS}(o, t^o, K) = \begin{array}{|l} s^{att} \\ \hline s^{att} : \text{ATT}(o, \langle \text{PERCEIVE}(K) \rangle, \langle \textit{now}, t^o \rangle) \\ \hline t^o \subseteq s^{att} \end{array}$$

The structure of the definition is reminiscent of that for operator MEM (13), since both introduce an attitude state, and the internal *now* of that attitude state is mapped to the location time of the attitude. However, if for MEM the holder of the attitude state is always i , i.e. the speaker or the holder of another attitude in whose scope

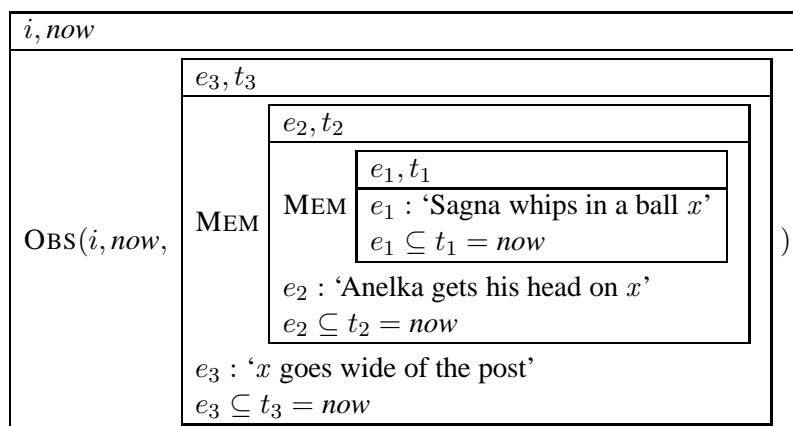


Figure 5: Iconic update. A DRS for (15)

MEM is interpreted, for OBS the perceiving individual o can (though need not) be distinct from i . The second difference concerns the update time: the time of the memory is always in the past, whereas the time of observation can be in the past, present, or future with respect to the actual *now*. In other words, MEM implements the idea that we can only remember our own past, whereas OBS implements a (shift of) perspective—seeing the world through one’s own or someone else’s eyes.

Reportive present: The most straightforward application of iconic update is to the interpretation of live reports, which feature the characteristic *reportive* use of the present tense. This is the case where the observer is the speaker, and the perception time coincides with the utterance time. Figure 5 shows the result of iconic update for the fragment (15) of a live commentary on the France–Uruguay match at the World Cup 2010.⁷

- (15) a. Sagna whips in a decent enough ball from the right...
 b. Anelka gets his head to it!
 c. But his glancing header goes wide of the post!

The difference to update from direct perception, is that the recursive embedding under the MEM operator does not apply to the top level DRS, but to the DRS K in the scope of the OBS operator. However, since the observer’s *now* is mapped to the discourse *now*, this does not have any effect on the temporal interpretation. Just as in the case of direct perception, it is always the last reported event that is happening *now*, whereas the previously reported events “move” into the past. The only real difference is that the observer is not the hearer. The top level DRS always represents the belief state of the recipient of the information, i.e. the observer in

⁷Published at: <http://uk.eurosport.yahoo.com/football/world-cup/uruguay-france-361810.html>. For reasons of space, various details are left out of the DRS in Figure 5.

the case of direct perception, and the hearer—the one who perceives the utterance events—in communication. OBS makes it explicit that the events are not perceived by the hearer, but by the speaker.

It is not difficult to see now how iconic update avoids the problem of unstable common ground, which arises with the classical update by communication when reporting changes in the world that take place while the discourse is in progress. Knowing how football is played, one can infer from the first utterance ‘Sagna whips in a ball from the right’ that the result state of that event holds, i.e. the ball is somewhere in the field, in some more or less central part of it. One could represent this by adding a state s_1 and corresponding conditions to the observer’s DRS:⁸

$$(16) \quad \text{OBS}(i, \text{now}, \left(\begin{array}{l} e_1, s_1, t_1 \\ e_1 : \text{‘Sagna whips in a ball } x\text{’} \\ e_1 \subseteq t_1 = \text{now} \\ t_1 \subseteq s_1 \\ s_1 : \text{‘ball } x \text{ is in the centre’} \end{array} \right))$$

Similarly, when we interpret the utterance (15c), which says that the ball (after having been played towards the goal) misses the goal, we can infer that it is somewhere left or right of the goal, i.e. definitely at the periphery of the field, if not out, and in any case not in the same place where it was after the throw in, so we could add the conditions: s_3 : ‘ball x is close to goal y ’ and $t_3 \subseteq s_3$ (and recall that $t_3 = \text{now}$, cf. Figure 5). Despite the fact that both s_1 and s_3 hold ‘now’ we do not infer that the ball is at two different locations simultaneously, because by the time we interpret (15c), the DRS representing (15a) is embedded under a MEM operator, so the two instances of *now* are ultimately mapped to two distinct time intervals in the top DRS (once we “unpack” the content of MEM, cf. Figure 4 in the previous section). If the two utterances were interpreted with respect to the same *now*, as it happens in non-iconic update by communication, we would derive a contradiction.

This is the explanation why the present tense does not go easily with events in its “normal” uses, while it does in its reportive uses. The normal uses are uses under non-iconic update, which, for good reasons, should not be employed for talking about events happening as the discourse is in progress. Reportive uses of the present tense always require iconic update. Next, we will see how the notion of iconic update can be applied to historical present and to past tense narrative.

⁸It might seem that condition $e_1 \subseteq t_1 = \text{now}$ is in contradiction with the condition $t_1 \subseteq s_1$ because they jointly entail that the event e_1 takes place while the result state s_1 holds, but intuitively, the result state only comes into being immediately after the event. However, this is a purely technical problem, which can be given a consistent solution, and on which we will not waste our time here. For the time being we could assume that the event only covers the initial time point of *now*, whereas the state covers the whole of *now* and possibly some time after *now*, but not the time before *now*. In this case the state and the event do overlap, but only in the initial moment of the state.

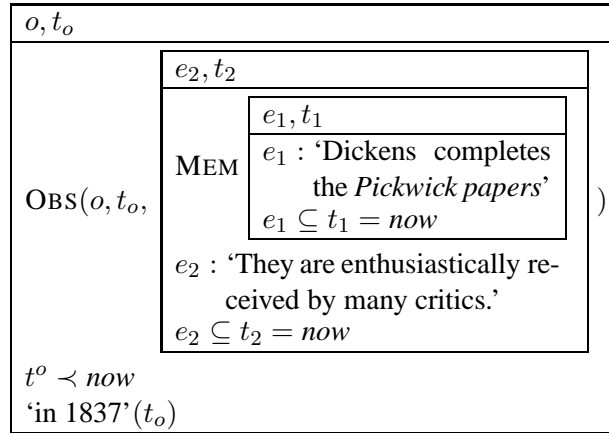


Figure 6: Historical present. Iconic update of (17a) with (17b).

Historical present: Historical present works just like reportive present, except that the observation time is located in the past and the observer may, but need not be identical with the speaker. The result of iconic update with the first two utterances of (17) (from Zucchi, 2005) is shown in Figure 6.

- (17) a. In 1837, Dickens completes the *Pickwick papers*.
 b. They are enthusiastically received by many critics.
 c. He moves to York
 d. and marries his grand-niece Joan.
 e. In 1838, they are divorced again.

The temporal modifier *in 1837* is an explicit indication that the present tense should be interpreted in a context shifted to the past. Notice that the modifier applies to the observation time t^o rather than to any of the event times.

There is a potential problem here. After update with the first utterance (17a) the *now* of the first event e_1 is in the immediate scope of the OBS operator, which means that it is mapped to the observation time t^o and is located in 1837. However, after update with the second utterance (17b), e_1 is not in the immediate scope of OBS anymore, but being in the scope of the MEM operator, lies in the past of the *now* of the second event, which is now mapped to t^o . Thus the inference that e_1 happens in 1837 is not granted anymore. Under iconic update, the previously mentioned events move literally into the past, and nothing is there to stop them from moving out of 1837 back into 1836. In other words, the system, as specified so far, is non-monotonic in a unintended way.

One can think of various ways to fix this problem. What we really want is that all the events in the scope of the same OBS are located within its observation time interval. One possible way to reach this is to stipulate explicitly that all the inferences from previous updates are preserved. Another possibility would be to make the MEM operator keep track of a “big now”—a variable that recursively

builds up a sum of all the times of events processed so far. It is then this big *NOW* that is anchored to the observation time t^o . The definitions of MEM and OBS would have to be modified as follows:

$$(18) \text{ MEM}(K) = \begin{array}{|l} t, T, s^{att} \\ \hline s^{att} : \text{ATT}(i, \langle \text{PERCEIVE}(K) \rangle, \langle \text{now}_1, t \rangle, \langle \text{NOW}_1, T \rangle) \\ t \subseteq s^{att} \\ t \prec \text{now}_2 \\ \text{NOW}_2 = T \oplus \text{now}_2 \end{array}$$

$$(19) \text{ OBS}(o, t^o, K) = \begin{array}{|l} s^{att} \\ \hline s^{att} : \text{ATT}(o, \langle \text{PERCEIVE}(K) \rangle, \langle \text{NOW}, t^o \rangle) \\ t^o \subseteq s^{att} \end{array}$$

By default, i.e. if not specified by a MEM operator, the big *NOW* is set equal to the ordinary *now*. The result is that t^o represents the sum $t_1 \oplus t_2 \oplus t_3 \oplus t_4$ of all the event times from (17a)–(17d), so *in 1837* modifies the whole sum: ‘in 1837’ ($t_1 \oplus t_2 \oplus t_3 \oplus t_4$).

This view also implies that the modifier *in 1838* in (17e) must open a new observer’s perspective. The observer can be the same, but the time of observation must be distinct from t^o located in 1837. This should reflect the intuition that temporal frame adverbials tend to introduce discontinuity into the flow of the narrative (see e.g. Roßdeutscher and von Stutterheim, 2006). If a sequence of events embedded under a single OBS operator could be compared to a continuous video recording, the transition to a new perspective is like a cut point in the film.

Simple past: Finally, we go back to our first example of past tense narrative, repeated in (20).

- (20) a. Josef turned around.
 b. The man pulled his gun from its holster.

The past tense locates the events before the discourse *now* ($t_1 \prec \text{now}$, $t_2 \prec \text{now}$), i.e. the *now* of the speaker and the hearer. However, in order to perform iconic update we need to get access to the *now* of the events themselves. One could introduce an optional rule that allows to represent an event in the past as an event *observed* in the past. This yields the representation in Figure 7 (using standard non-iconic update). The two occurrences of the past tense introduce two past perspectives whose observers (o_1, o_2) and observation times (t_1^o, t_2^o) are *a priori* distinct. However, we will assume that as long as nothing forces us to interpret them as truly distinct (such as the conflicting temporal locations *in 1837* vs. *in 1838*, as in the previous example), they are interpreted as identical. That is, by default $o_1 = o_2$ and $t_1^o = t_2^o$. This can be seen as a consequence of applying the general pragmatic principle of topic continuity (Givón, 1983; Jasinskaja, 2007), *NEW (Blutner and Zeevat, 2003), or *Maximise Discourse Coherence* (Asher and Lascarides, 2003),

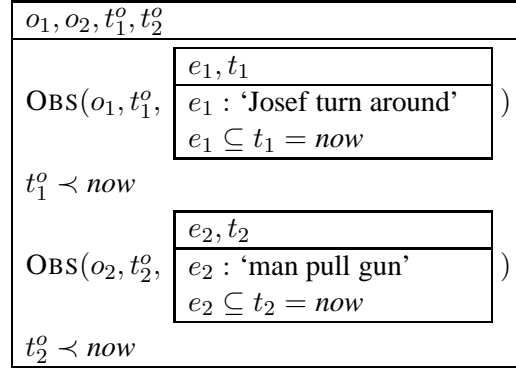


Figure 7: Simple past. Representation of (20) using a past perspective: ‘event e_1 was observed to happen in the past’.

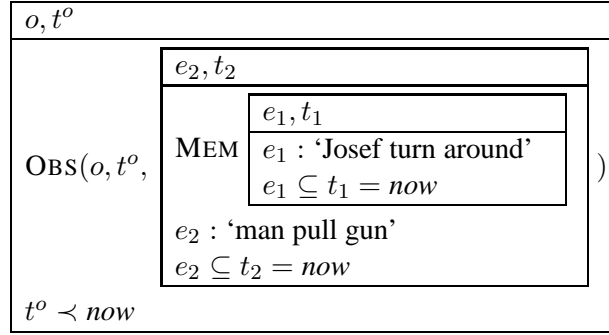


Figure 8: Simple past. Iconic update of within the scope of OBS.

which advises against introducing new objects into the universe of discourse and bids one to assume that reference is made to previously introduced objects as far as possible. But once the observer and the observation time are the same, nothing prevents us from putting both event representations within the scope of the same OBS operator. Of course, within the scope of OBS, the DRSs should be combined in the iconic fashion, since we are now dealing with observations, rather than utterances, cf. the ‘iconic conjunction’ \wedge_i in (21) and its definition in terms of iconic update in (22).

$$(21) \text{ OBS}(o, t, K_1) \wedge \text{ OBS}(o, t, K_2) \text{ iff } \text{ OBS}(o, t, K_1 \wedge_i K_2)$$

$$(22) K_1 \wedge_i K_2 \text{ iff } \text{ MEM}(K_1) \wedge K_2$$

The result DRS is shown in Figure 8. Notice that it has the same structure as the DRS in Figure 6 for the example in historical present. The temporal sequence of reported events is established by recursive application of the MEM operator, whereas the whole sequence is located in the past via anchoring to a past observation time.

To summarise, in this section it was shown how the standard predictions of the

DRT account of temporal progression in narrative discourse can be reconstructed in a theory based on an explicit notion of resemblance between the process of discourse interpretation and the process of direct perception of events in the world. An account of reportive and historical uses of the present tense came for free along with an account of past tense narrative. But all these predictions were previously obtained without an explicit representation of past memories and observer’s perspectives, within technically simpler theories, such as Kamp and Reyle’s (1993) theory based on reference and perspective times, as well as Zucchi’s (2005) account of reportive and historical present. Does the present proposal gain any new empirical grounds as compared to those theories? The next section gives a brief example of a case in point.

4 Beyond time

Do we really need an explicit representation of the observer in our model, or can we reach a perfectly adequate discourse interpretation just by directly manipulating aspects of temporal reference? If perspective taking in discourse only affected its temporal dimension, then the latter would perhaps be a simpler option. However, perspectivisation does not only concern time. As we know from literary studies as well as from studies at the interface of literature and linguistics (e.g. Banfield, 1982; Fludernik, 1993; Ryan, 2004), the holder of the perspective, even if it is a fictitious individual neither identical with the speaker nor with any of the characters in the story, can be attributed a variety of properties, including beliefs, attitudes, and even, to some extent, action. If our discourse model is to represent these inferences, then the possibility of explicit reference to the observer is indeed required.

Just to get a glimpse of where this approach can take us, consider the following pair of examples:

- (23) a. John entered the room. (*e*)
 b. Max was sitting in his chair. (*s*)
- (24) a. Max was sitting in his chair. (*s*)
 b. John entered the room. (*e*)

One has a clear intuition that in (23) the observer is moving along with John: Max sitting in his chair is the view that opens as the observer “enters” the room. In (24) the observer is in the room with Max and sees John come in.

Until now we have only represented the belief or the perception *state* of the observer that results from consecutive update with new information. This state can be thought of as the result state of a belief change or perception *event*, i.e. an update event. If such update events are also represented explicitly in the DRS, the difference between (23) and (24) can be easily expressed. The DRS in figure 9 shows this for (23). In addition to the observation state s_o which accumulates all the information acquired by perception during the period t^o , the DRS also represents the observation state s_2^o which corresponds to the perception of just the state of

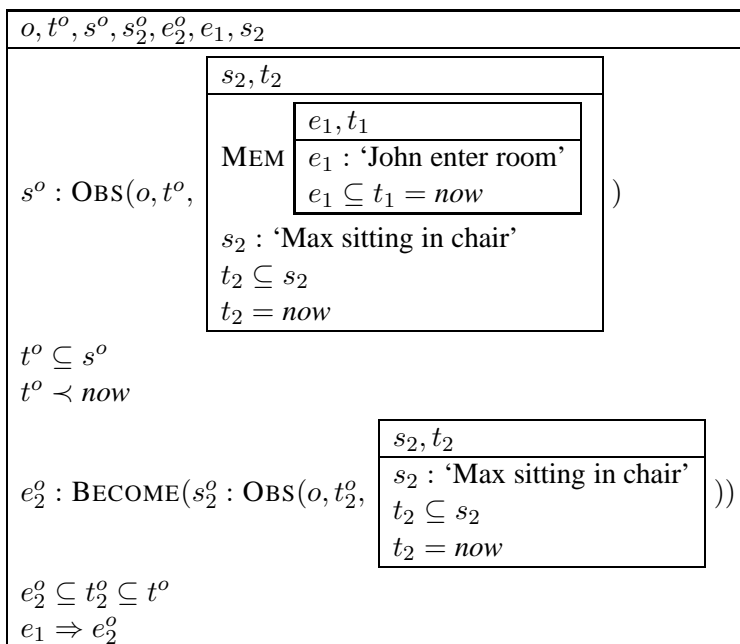


Figure 9: DRS for (23)

affairs described by the second sentence ‘Max was sitting in his chair’ and the event e_2^o that brings about that state. That is, e_2^o is the event of update of the previous observation state with this new perceptual input.⁹

A crucial step towards capturing the intuition behind (23) is to establish a causal relation between the event e_1 of John entering the room and the update event e_2^o of the observer perceiving that Max is sitting in his chair, represented as $e_1 \Rightarrow e_2^o$ in figure 9. How exactly this inference comes about will not be spelled out here. We assume that its nature is essentially the same as that of the causal inference in ‘John fell. He broke his leg.’ The pragmatic principle of coherence maximisation (such as the MDC in SDRT, Asher and Lascarides, 2003) prefers interpretations that make the semantic links between sentences as tight as possible. In particular, it makes the hearer infer causal relations on top of mere temporal succession wherever possible. It is natural to assume that such causal inferences can affect not only the actual reported events and states, but also the update events and observer’s perspective states.

The event of John entering the room can cause the observer to perceive Max sitting in his chair under two most plausible scenarios: either John is the observer,

⁹The notation is somewhat changed as compared to that in previous sections. The top DRS in figure 9 contains an explicit referent s^o for the observer’s perspective state, which used to be implicit in the OBS operator, cf. s^{att} in (19). The objects e_1 and s_2 in the universe of the top DRS are the external anchors of e_1 and s_2 of the respective embedded DRSs, i.e. they represent the event e_1 and the state s_2 in the perspective of the hearer (rather than the observer).

or John entering the room causes a distinct observer (“John’s shadow”) to follow him, which in turn causes the observer to perceive Max sitting in his chair. In both cases, we infer that the observer enters the room.

Obviously, in (24) this inference does not go through because the observer sees Max sitting in his chair before John enters the room. Since causes precede their effects, the event of John entering the room cannot be the cause of the update event.

This brief example shows what can be gained by the explicit representation of the observer. Inferences of this type cannot be modelled if perspective taking in discourse is reduced to temporal perspective alone.

5 Conclusions and outlook

This paper does not do much more than present an idea of how the notion of iconicity, in particular, temporal iconicity in narrative discourse, can be made explicit in a formal semantic framework, such as DRT. The central idea is that the update time *now* is processed differently in (non-narrative) discourse and in direct perception, and that the interpretation of narrative discourse deviates from the normal procedure for *now* in communication, miming direct perception. The proposed approach models chronological interpretation without making use of reference times or presuppositions of a temporal relation to previously mentioned events, as is standard in classical DRT. It covers the standard cases of past tense narrative, whereas an account of reportive and historical present comes along for free. In the vein of SDRT, the present approach presupposes that the fact that we have to do with a narrative is established on independent grounds, however instead of simply stipulating that events in a *Narration* follow one another, it models temporal sequentiality as a consequence of “vicarious experiencing of events in the real world”, by explicitly representing the reported events as being perceived by a real or fictitious observer. As long as we are only concerned with temporal perspective, this might seem like an unnecessary complication of the theory. However, if our theory is to model other inferences involving the observer, such as the location and motion of the observer in space, as in the example we discussed, or the observer’s beliefs and attitudes, then a representation of the kind we proposed becomes indeed necessary.

There are many more open questions that this approach gives rise to. First of all, the formal details still need to be worked out. There is a major caveat in DRT (but also in any other dynamic semantic framework I am aware of) which makes it, as is, not entirely suitable for the task. DRT only provides explicit representation for belief *states*, but not for update *events*. Update in dynamic semantics is conceived of as a function from belief states to belief states, but not as an event that happens at a particular time and place, involves certain participants, etc., that is, essentially the same kind of thing as an event of falling, pushing, pulling one’s gun from its holster, etc. Such a representation is necessary in order to build a perfect analogy between update by direct perception and update in discourse, be that iconic or non-iconic update. A proper solution to this problem will preempt

most other questions related to semantics that the present proposal might rise. For instance, the sortal difference between the object of belief (standardly, proposition or fact) and the object of perception (standardly, event or state) was completely ignored in this paper. If we had a systematic way of representing the relationship between beliefs as (sets of) propositions, on the one hand, and belief *states*, which function as prestates and result states of update events, on the other hand, the analogy developed in this paper would not require such a blunt confusion of sorts.

Also the empirical consequences of the present proposal still need to be properly explored. In the domain of temporality, our goal was to reconstruct the standard predictions of existing theories, however in at least one point the proposed approach deviates from the mainstream. Standardly, it is events that push the narrative forward, whereas states do not. In our approach, it is neither events, nor states, but the sequentiality of perception, which applies to events and states equally. If according to the standard approach, the discourses in (25) receive the same interpretation: the two states overlap each other (and some previously mentioned event) regardless of the order, in our approach (25a) and (25b) are not equivalent, as they imply different orders of perception by an observer. Both positive and negative consequences of this theoretical turn still need to be properly studied.

- (25) a. Mary was making coffee. Sue was talking with a neighbour.
b. Sue was talking with a neighbour. Mary was making coffee.

As far as implications beyond temporality are concerned, it is potentially a vast field which only starts to be explored by formal linguists. The relevant phenomena include various kinds of perspective shifting that make the daily bread of literary narratology, including *voice* (Who is speaking? What kinds of narrators should be distinguished?: Genette, 1986, 1988; Fludernik, 1993), *mode* (Which perception point defines the field of vision in the narrative? Whose perspective orients the text?: Genette, 1986, 1988; Ryan, 2004), and *speech/thought representation* (Which epistemic background underlies the narrative? Which protagonists can reasonably hold this epistemic background?: Banfield, 1982; Fludernik, 1993; Pascal, 1977). In linguistics, Apresian (1986) and Paducheva (1996) were among those who argued for explicit representation of the observer in the analysis of a whole range of linguistic phenomena. All these phenomena are obvious candidates for treatment within the approach proposed in this paper, but the work still needs to be done.

References

- Apresian, J. D. (1986). Integral'noe opisanie jazyka i tolkovyj slovar'. *Voprosy Jazykoznanija*, 2:57–70.
- Asher, N. and Lascarides, A. (2003). *Logics of Conversation*. Studies in Natural Language Processing. Cambridge University Press.

- Banfield, A. (1982). *Unspeakable Sentences: Narration and Representation in the Language of Fiction*. Routledge, Kegan Paul, London.
- Bennett, B., Hoffman, D., and Prakash, C. (1989). *Observer mechanics: A formal theory of perception*. Academic Press.
- Blutner, R. and Zeevat, H., editors (2003). *Optimality Theory and Pragmatics*. Palgrave Macmillan.
- Dowty, D. R. (1986). The effects of aspectual class on the temporal structure of discourse: semantics or pragmatics? *Linguistics and Philosophy*, 9:37–61.
- Fludernik, M. (1993). *The Fictions of Language and the Languages of Fiction*. Routledge, London.
- Genette, G. (1986). *Narrative discourse: An Essay in Method*. Blackwell, Oxford. Translated by Jane E. Lewin with a foreword by Jonathan Culler.
- Genette, G. (1988). *Narrative discourse revisited*. Cornell University Press, Ithaca, NY. Translated by Jane E. Lewin.
- Givón, T. (1983). Topic continuity in spoken English. In Givón, T., editor, *Topic Continuity in Discourse*, pages 347–363. John Benjamins, Amsterdam.
- Hinrichs, E. (1986). Temporal anaphora in discourses of English. *Linguistics and Philosophy*, 9:63–82.
- Jakobson, R. (1971). Quest for the essence of language. In *Selected writings, Vol. 2: Word and language*, pages 345–359. Mouton, The Hague, Netherlands. First published in 1965.
- Jasinskaja, E. (2007). *Pragmatics and Prosody of Implicit Discourse Relations: The Case of Restatement*. PhD thesis, University of Tübingen.
- Jasinskaja, K. and Karagjosova, E. (2011). Elaboration and explanation. Paper presented at Constraints in Discourse 4. URL: .
- Kamp, H. and Reyle, U. (1993). *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer Academic Publishers.
- Kamp, H., van Genabith, J., and Reyle, U. (2005). Discourse representation theory. To appear in *Handbook of Philosophical Logic*.
- Nöth, W. (2001). Semiotic foundations of iconicity in language and literature. In Fischer, O. and Nänny, M., editors, *The Motivated Sign: Iconicity in Language and Literature*, pages 17–28. John Benjamins.
- Paducheva, E. V. (1996). *Semantičeskie issledovanija. Semantika vremeni i vida v russkom jazyke. Semantika narrativa*. Moskva.
- Papafragou, A., Hulbert, J., and Trueswell, J. (2008). Does language guide event perception? Evidence from eye movements. *Cognition*, 108(1):155–184.
- Partee, B. H. (1984). Nominal and temporal anaphora. *Linguistics and Philosophy*, 7(3):243–286.
- Pascal, R. (1977). *The Dual Voice: Free Indirect Speech and its Functioning in the Nineteenth-Century European Novel*. Manchester University Press, Manchester.
- Reyle, U., Rossdeutscher, A., and Kamp, H. (2007). Ups and downs in the theory of temporal reference. *Linguistics and Philosophy*, 30(5):565–635.
- Roßdeutscher, A. and von Stutterheim, C. (2006). Semantische und pragmatische

- Prinzipien der Positionierung von *dann*. *Linguistische Berichte*, pages 29–60.
- Ryan, M.-L. (2004). Introduction. In Ryan, M.-L., editor, *Narrative Across Media: The Languages of Storytelling*. University of Nebraska Press, Lincoln.
- Schmiedtová, B. and Sahonenko, N. (2008). Die Rolle des grammatischen Aspekts in der Ereignis-Enkodierung: Ein Vergleich zwischen tschechischen und russischen Lernern des Deutschen. In Walter, M. and Grommes, P., editors, *Fortgeschrittene Lernervarietäten: Korpuslinguistik und Zweitspracherwerbsforschung*, *Linguistische Arbeiten*, pages 45–71. Niemeyer.
- Schmiedtová, B., von Stutterheim, C., and Carroll, M. (to appear). Language-specific patterns in event construal of advanced L2 speakers. In XXX. XXX.
- Segal, E. (1995). Cognitive-phenomenological theory of fictional narrative. In Duchan, J., Bruder, G., and Hewitt, L., editors, *Deixis in narrative: a cognitive science perspective*, pages 61–78. Erlbaum, Hillsdale, NJ.
- Smith, C. (1991). *The Parameter of Aspect*. Kluwer, Dordrecht.
- Zucchi, S. (2005). The present mode. In Carlson, G. N. and Pelletier, F. J., editors, *Reference and quantification: The Partee effect*, pages 1–28. CSLI Publications, Stanford.
- Zwaan, R. (1999). Situation models: The mental leap into imagined worlds. *Current Directions in Psychological Science*, 8(1):15–18.
- Zwaan, R., Madden, C., and Stanfield, R. (2001). Time in narrative comprehension. In Schram, D. H. and Steen, G., editors, *The psychology and sociology of literature: In honor of Elrud Ibsch*, pages 71–86. John Benjamins.