There is no Opposition between Formal and Cognitive Semantics

FRITZ HAMM, HANS KAMP, and MICHIEL VAN LAMBALGEN

1. Introduction

The history of modern semantics is characterised by two research traditions which are based on radically different views concerning both conceptual motivation and the purpose of semantic research.

Realistic semantics conceives of semantics as characterising the relationsship between linguistic expressions and reality. In most cases this relationship is explicated by means of modeltheoretic concepts. The following quote from one of the founding fathers of realistic semantics clearly rejects a mentalist stance.

> I distinguish two topics: first, the description of possible languages or grammars as abstract semantic systems whereby symbols are associated with aspects of the world; and second, the description of the psychological and sociological facts whereby a particular one of these abstract semantic systems is the one used by a person or a population. Only confusion comes of mixing these two topics.
> Lewis [1972], p 170

From Lewis' quote above it is quite clear what realistic semantics does not aim at. It is not about psychological or sociological facts concerning the use of a linguistic system but about the relationship of an abstract linguistic system to aspects of the world. The conclusion is that the relationship between language considered as an abstract system and the world is at least in principle independent of meanings grasped by a mind.

A meaning for a sentence for instance is something which determines under what conditions the sentence is true or false. The meaning determines how the truth conditions, the extension, depends on relevant factors such as facts about the world, on the speaker, on the surrounding

Theoretical Linguistics 32–1 (2006), 1–40 0301–4428/06/0032–0001

discourse etc.. The collection of these relevant factors is called an *index* by Lewis and the sort of things which determine how something depends on something else are functions in the set theoretical sense. An intension for a sentence is any function from indices to truth values and an intension for names is any function fom indices to things. This idea is generalised to cover other syntactic categories as well. Therefore the basic notions of realistic semantics are reference, truth and, based upon these concepts, inference[1].

Cognitive oriented research in semantics considers the investigation of the relationship between natural language expressions and mental structures as the major topic of semantic research.

For example, Jackendoff – a semantisist working in the mentalist tradition – doubts the foundational role of the above concepts and claims that they are derivative on conceptual structure. He describes the aims of his own conceptual semantics in the following way:

> Conceptual Semantics ... is concerned most directly with the form of the internal mental representations that constitute conceptual structure and with the formal relations between this level and other levels of representation.... Conceptual Semantics is thus a prerequisite to [truth functional] semantics: the first thing one must know about an English sentence is its translation into conceptual structure. Its truth conditions should then follow from its conceptual structure plus rules of inference, which are stated as well in terms of conceptual structure.
> Jackendoff [1994], p 132

Conceptual structures are generated – similarily to phrase structures by phrase strtucture grammars in syntax – by an algorithmic dervice[2]. Therefore semantics constitutes a generative domain of its own, independent of syntax to which it is linked by so called correspodence rules. One may think of the formal representations of conceptual structures as labeled bracketings where the labels are drawn from the major conceptual categories such as *event*, *thing* and *path*.

---

[1] Although this extremely short description does not justice to the subtleties of more modern versions of realistic semantics Lewis' account of the conceptual foundations is still valid and unsurpassed in its clarity.

[2] In Foster [1992] Carol Foster argues that the notion of *algorithm* is central for cognitive science since it introduces such important concepts as *processing* and *complexity*.

In his reply to Abbott Jackendoff [1998] summarises his position as follows:

> . . . : I am interested in studying the properties of the human mind, and I think that this is ultimately a more productive context for examining human language.
> Jackendoff [1998], p 211

We wholeheartedly agree but nevertheless insist that this position too requires explicit formalisation in order to ensure the computability which is fundamental to cognition. It is the purpose of this paper to motivate and outline a system which achieves this. The last section of this paper illustrates the formal devices of a semantic system based upon the notion of computation.

We start with discourse representation theory (DRT) since this was historically the first system which could address both questions arising in a realistic and in a cognitive framework[3].

2. Discourse Representation Theory

2.1. Conceptual motivation

There exist a number of introductions to DRT in the literature, some of which have been around for a good many years. So we do not make it a task of this paper to reiterate what has been explained in extenso elsewhere. Instead we concentrate on some of the conceptual motivations behind DRT.

The principal respect in which DRT differs from the formal approaches to the analysis of meaning in natural language that existed at the time when it was conceived is the attention it pays to the systematic ways in which the interpretation of words and sentential constructions may depend on the discourse context, such as it is given by the sentence or sentences with which the given sentence co-occurs in a connected discourse or text, and to the intersentential semantic relations that are created by such dependencies. The type of interaction for which DRT and the more

---

[3] Moreover Jackendoff [2002] uses a semiformal notational variant of DRT for his notion of *referential tier*.

or less simultaneously developed File Change Semantics of Heim[4] became known for was the solution they offered to the so-called donkey pronoun problem exemplified by (1).

(1) If I rent a car, I only use it to visit my relatives.

In the case of DRT, however, the original impetus for this approach came from a certain view of the semantics of tense. The actual starting point was an attempt to account for the differences between the Passé Simple and the Imparfait. One difference between these two tenses concerns the way in which sentences containing them are semantically connected with the sentences that precede them. A classical example is (2), where the difference between PS and Imp in the second sentence is crucial to the meaning that is conveyed by the two sentences together.

(2) Quand Alain ouvrit les yeux, il s'aperçut sa femme qui était debout près de son lit.
When Alain opened his eyes, he saw his wife who was standing next to his bed.
a. Elle lui souriait.
She was smiling at him.
b. Elle lui souria.
She smiled at him.

As this contrast is described in DRT, the Imparfait in (2-a) (and similarly the past progressive in English) is interpreted as expressing a state or process that was going on while Alain opened his eyes and noticed his wife. In contrast, the PS sentence (2-b) is interpreted as describing an event which happened after the event described in the first sentence of (2), and presumably in response to that event. This example shows how the tense of a sentence constrains the way in which the event or state of the new sentence is temporally connected with that of the preceding sentence.

---

[4] See Heim [1982]. The original versions of FCS and DRT were equivalent in the predictions they make within the target domain which they share, and for this reason they are now often treated as a single theory. But as far as we can tell, there may nevertheless be certain differences between FCS and DRT as regards underlying conceptions and motivation, so what follows is intended as relating more narrowly to DRT (whether in the form presented originally in Kamp [1981] and then in greater detail in Kamp and Reyle [1993] or in the revised version of von Genabith et al. [2005]).

To formulate principles which capture the reference (and thus the semantic contribution) of tenses and other such expressions, requires essential reference to the discourse context and more often than not this is the 'global' context provided by the preceding sentences.

In order to get such a theory of sentence-transcendent interpretation actually to work in detail, it is necessary to make certain assumptions about the structure of the discourse context K_n resulting from the interpretation of a discourse segment $\langle S_1, \ldots, S_n \rangle$ which is needed for the interpretation of the immediately following sentence S_{n+1}. The use of K_n for this purpose imposes constraints on its structure that go beyond the requirement that it correctly represent the truth conditions of $\langle S_1, \ldots, S_n \rangle$. To capture these additional constraints DRT proposed that the context representations which result from discourse interpretation (its so-called Discourse Representation Structures, or 'DRS's) not only determine the truth conditions of the interpreted discourse segment, but also are distinguished by certain formal properties which are not essential to the represented truth conditions but become important when they are used as contexts.

This apparent need for context representations with formal properties that go beyond the represented truth conditions can be seen as circumstantial evidence for the hypothesis that natural language interpretation involves a level of semantic representation whose representations reflect aspects of the represented input that cannot be recovered from its truth conditional content. That this level cannot be a level of syntactic representation is implied for one thing by the fact that its representations are in general representations of multi-sentence discourse, and not just of single sentences, on the assumption that syntactic structure is limited to single sentences and their constituents. A further argument for the view that DRSs constitute a distinct level of genuinely semantic representation follows below.

2.2. DRT and the cognitive dimension of language use

DRT's claim that there are features of natural language the analysis of which requires a distinct level of discourse representation is consonant with a cognitive perspective on the nature of natural language meaning: Meaning in natural language manifests itself as the semantic competence of the language user, this competence is demonstrated in the interpretation and production of utterances, and language production and interpretation involve mental representations, which are derived from linguistic input in language interpretation and converted into linguistic output in language production. For someone who thinks of meaning along these lines it is tempting to see the formal properties of discourse contexts which DRT identifies as defining the interpretational possibilities for anaphoric pronouns as features of the mental representations which are constructed in the course of interpreting a text or piece of discourse; and this encourages a view of DRSs as models for mental representations, which capture some of the formal properties of those representations in addition to their truth conditional content.

Although this cognitive perspective was one of the conceptual motivations behind the development of DRT, it was downplayed in DRT's early years out of the concern that this might detract from the theory's potential as a form of formal semantics. Even as it was, DRT was soon criticized for its representational position, in particular by Groenendijk and Stokhof[5], who proposed an (almost) representation-free alternative. However, when other phenomena are considered, besides those treated in the early presentations of DRT, the need for some mode of representation becomes more prominent. A salient case are plural pronouns. As argued at some length in Kamp and Reyle [1993], plural pronouns often have anaphoric antecedents which must be constructed from the "raw material" that discourse contexts make available. The rules which are needed for the construction of such antecedents – such as summation and abstraction – cover between them many of the sets whose existence can be derived from the truth conditional content of the discourse context. But interestingly they do not cover the full range of those sets. Moreover, the particular part that they do cover reflects an interpretation regime that is specific to plural pronouns and does not extend to other types of definite noun phrases, such as definite descriptions.

Further complications arise for so-called dependent pronouns (both plural and singular) that are found in sentences which follow quantificational statements, as in examples like (3), first discussed by Hintikka

---

[5] Groenendijk and Stokhof [1991], Groenendijk and Stokhof [1990].

(3) I gave a present to each of the children in the orphanage. Most children opened them/it right away.

Like cases of plural pronoun anaphora which involve set abstraction, the interpretation of the second sentence of (3) requires knowledge of the quantificational structure of the preceding sentence or sentences. In DRT the interpretation of plural and dependent pronouns is treated, like that of non-dependent singular pronouns, along representational lines. For instance, cases like (3) can be dealt with because quantificational structure is explicitly represented, and thus available as part of the context representation when a pronoun involving this kind of anaphora is up for interpretation. Here too there have been alternative proposals that follow the spirit of Dynamic Semantics. In these proposals the needed quantificational structures are not coded as components of context representations, but instead the quantificational dependencies are treated as additional structure of the situations which the discourse segments describe.[6]

Whether the approach these proposals exemplify should be preferred to treatments in the style of DRT is a question we leave to others. What matters here is this: Wherever the theory of one's preference locates the extra information that is needed to account for the interpretational options of dependent pronouns, it is information that is available only when the discourse has made it available. That it should be technically possible to treat such information as additional structure of the denotations of discourse segments in actual and possible situations or worlds isn't all that surprising. But the effect of doing this is that denotations are made to incorporate aspects of how the given discourse describes the situation or situations it targets, and not just what it says about them. To think that representationalism could be eliminated just by relocating information that is contributed by the describing discourse in this manner would clearly be an illusion.

Our arguments for the thesis that some sort of representational structure of discourse contexts is needed have been of a purely 'functional'

---

[6] See in particular van den Berg [1996], Elworthy [1992], Elworthy [1995], Krifka [1996] and Nouwen [2003]. To be more precise, these proposals impose the extra structure on the denotations of the relevant discourse segments in the different relevant models, where the models are playing, as usual in model-theoretic semantics, the part of the situations (or possible worlds) that the discourse describes or could have described.

nature: Unless, the arguments went, certain structural features of the antecedent discourse context are taken into account, many types of discourse anaphora and the inter-sentential meaning relations which result from them cannot be explained. None of these arguments appeal to cognitive considerations, according to which interpreting and producing language must involve mental representations. Of course this doesn't mean that they would have less relevance for someone who believes in mental representations than they are for those who see natural language semantics as an enterprise that need not and should not make commitments in this direction. For one thing, they matter insofar as they indicate that mental representations cannot be purely syntactic in a sense of the term that is consistent with current conceptions of syntax.

Once a semantic representation has been obtained it will be normally exploited in further information processing. Many of these processes (and perhaps all) exploit semantic representations as premises for various kinds of inference, either on their own or, more usually, in combination with others, some or all of which may be of non-linguistic origin (e.g. visual perception). The repertoire of inference principles that are used by human information processors is still poorly understood. It is our conviction, however, that inference principles of classical logic are part of this repertoire since they play an important role in non–standard systems too.[7] If this is so, then mental representations that are used as premises in human reasoning should be of a form that is accessible to such inference principles. This is one reason for insisting that semantic representations have such a form; and thus, if one thinks of DRSs as modelling the way in which the mind represents meaning, it is natural to require of them that they display such a form as well. In the case of DRT the consequence relation is formally specified as follows: DRS K_2 is a logical consequence of DRS K_1 iff every model of K_1 is also a model of $K_1 \oplus K_2$. A system of inference rules that axiomatises this relation for first order DRS languages and which is adapted to the special feature of DRSs can be found in Kamp and Reyle [1996]. The system is adapted to the special features of DRSs but is at the same time close to familiar inference systems for first order predicate logic.

---

[7] See the study Stenning and van Lambalgen [2005] for concrete examples.

The claim that classical inference principles are included in the repertoire of the human cognitive system does not mean that the entailment relations implemented in human reasoning processes must coincide with the classical consequence relation. On the contrary, it is widely held – and we concur – that as a rule these entailment relations are stronger. This is so in particular when people reason about temporal and causal relations between events and states of affairs and more particularly yet when they reason about the contents of temporal discourse, such as narrative descriptions of shorter or longer episodes. In (2) we encountered two brief and comparatively simple examples of such descriptions and in the following sections we will discuss a number of others. In fact, episode descriptions provide a good illustration of another point, viz. that inferencing not only comes into play after semantic representations have been constructed, but is needed also during the construction of these representations (Rossdeutscher and Reyle [2000], Reyle et al. [2005]). The relevant entailment relations at the level of representation construction are often based on the assumption that the situation about which one reasons consists only of those entities whose existence is entailed by the context representation.

The interpretation of temporal discourse and the subsequent exploitation of the information that is thereby obtained thus present a dual challenge. On the one hand there is the problem of accounting for how the different constituents of the sentences composing such a discourse make their interacting contributions to its semantic representation. On the other both the construction of such representations and their subsequent inferential use require modes of inference that have now been identified in general terms but are much in need of further clarification. It is this double challenge that the marriage between CLP (Constraint Logic Programming) and DRT is intended to meet.

3. Tense, aspect and all that

It is the purpose of verb tenses and (lexical and grammatical) aspect to generate the order and structure of the events described by a piece o discourse. The implication here is that it is not very useful to study tense and aspect at the sentence level, as generative approaches to linguistics maintain; tense and aspect really come into their own only at the discourse level. Examples will be supplied below. The book van Lambalgen and Hamm [2004] attempts to justify the assumption that only by looking at the cognitive construction of time will we be able to understand how time is encoded in linguistic constructions – with the added bonus that predictions about cognitive processing of tense and aspect can be derived.

We believe that formal semantics must be relevant to explaining language comprehension and production, over and above getting the linguistic data normally taken into account in formal semantics (truth conditions of sentences and entailments between sentences in context) right. This means that one is not completely free to choose a formalism, subject only to the constraint of consistency with the data; after all, some formalisms may not be 'executable' by the brain, e.g. because they are not computable at all, or if they are, because of limitations of working memory.

3.1. Tense and aspect in discourse

We take it as the essential purpose of tense and aspect to facilitate the computation of the structure of the events described in a narrative. We write 'facilitate', because tense and aspect cannot by themselves completely determine event structure. The following examples (4) will make clear what we have in mind: these feature mini-discourses in French all consisting of one sentence in the Imparfait and one in the Passé Simple. The structure of the set of events differs in each case, however.

(4) a. Il faisait chaud. Jean ôta sa veste. (Imp, PS)
It was hot. Jean took off his sweater.
b. Jean attrapa une contravention. Il roulait trop vite. (PS, Imp)
Jean got a ticket. He was driving too fast.
c. Jean appuya sur l'interrupteur. La lumière l'éblouissait. (PS, Imp)
Jean pushed the button. The light blinded him.

In the first case, the Imp-sentence describes the background against which the event described by the PS-sentence occurs. In the second case, the event described by the PS terminates the event described by the Imp, whereas in the third case the relation is rather one of initiation. These examples also show that world-knowledge in the form of knowledge of causal relationships is an essential ingredient in determining event structure. This knowledge is mostly applied automatically in computing event structure, but may be consciously recruited if the automatic processing leaves the event structure still underdetermined. It is the task of cognitive science to determine how this algorithm is actually implemented. (For some suggestions in this direction see section 3.8)

We hypothesize that there is an intimate connection between the ability to use tensed language and the general human capacity to form and execute plans. In its simplest form a plan consists of a sequence of actions – together with the times at which they have to be executed – which achieves a goal; but more complex plans are possible which also involve overlapping actions, such as for example drinking while walking. Part of this hypothesis is to see statements about the future, and especially those which are relevant to the interpreter's own future, as paradigmatic for what goes on in language comprehension generally. Interpreting such a sentence (and accepting its information as correct) amounts to the interpreter adjusting his model of the future in such a way that the sentence is true in it. The link between planning and linguistic processing is thus provided by the notion of goal: *we view a sentence* S *as a goal ('make* S *true') to be achieved by updating the discourse model.* Moreover, adjustment of the model will often have features reminiscent of planning in that the interpreter will adopt, as part of his modified model, assumptions about what will lead to the future state or event of which the sentence speaks.

The link between planning and statements about the present or past is arguably less direct. But here too we see a connection. In this case the connection involves not so much – or at any rate not only – the formation of plans, but their execution. Executing a plan involves keeping track of the successive actions of which it consists and to take note of those actions that have already been performed, seeing them as that part of the plan which has been dealt with; but it is also linked to the still future goal, with the agent's current now as juncture. In some cases where statements about the past or present are relevant to the future, and especially to the interpreter's own concern, understanding and accepting the sentence will have a similar effect on the interpreter's model of his world as processing statements about the future. In general, model adjustment for past tense (or present tense) need not have much of a direct impact on the interpreter's idea of what will happen to him. But in such cases too model adjustment takes essentially the same form.

3.2. Planning, causality and the ordering of events

In section 3.1 we formulated our main hypothesis as

> the ability to automatically derive the discourse model determined by a narrative (in conjunction with tacit world-knowledge) would have been impossible without the ability to compute plans to achieve a given goal.

In this section we provide a preliminary discussion of this hypothesis, as a preparation for the discussion of the formalism in section 3.4. The hypothesis will be seen to have two components: (i) planning subserves the construction of discourse models, and (ii) the human cognitive construction of time is built on our planning capacity, and tense and aspect systems reflect cognitive time, so that tense and aspect ultimately reflect features of planning. We will now discuss these components in turn.

(i) By definition, planning consists in the construction of a *sequence* of actions which will achieve a given goal, taking into account properties of the world and the agent, and also events that might occur in the world. The relevant properties include stable causal relationships obtaining in the world, and also what might be termed 'inertia', in analogy with Newton's first law. If a property has been caused to hold by the occurrence of an event, we expect that the property persists until it is terminated by another event. This is the inertial aspect of causality: a property does not cease to hold (or come to hold) spontaneously, without identifiable cause. Such inertia is a prerequisite for successful action in the world; and we will have to find a formal way to express it. It does however not suffice for successful planning.

Consider again the characterization of planning as setting a goal and devising a *sequence* of actions that will achieve that goal, taking into account events in, and properties of the world and the agent. In this description, 'will achieve' definitely cannot mean: '*provably* achieves in classical logic', because of the notorious frame problem: it is impossible to take into account all eventualities whose occurrence might be relevant to the success of the plan, but classical logic forces one to consider all models of the premisses, including those that contain farfetched possibilities. Therefore the question arises: how to characterize formally what makes a good plan?

A reasonable informal suggestion is: the plan works to the best of one's present knowledge. More formally, this idea can be reformulated semantically as: the plan achieves the goal in a 'minimal model' of reality; where a minimal model is characterized by the property that, very roughly speaking, every proposition is false which you have no reason to assume to be true. In particular, in the minimal model no events occur which are not forced to occur by the data, and only explicitly mentioned causal influences are represented in the model. This makes planning a form of nonmonotonic reasoning: the fact that

'goal G can be achieved in circumstances C'

does not imply

'goal G can be achieved in circumstances $C + D$'

The first claim of this paper (and of van Lambalgen and Hamm [2004]) can now be formulated as: planning computations underlie the construction of discourse models, which are in fact minimal models in the sense defined above.

(ii) This is the argument of the first part of van Lambalgen and Hamm [2004], and it is impossible to reproduce that argument in any detail here. Its brief outlines are as follows.

Physical time is not perceived in any literal sense; so cognitive time cannot be a direct reflection of perception of time; rather it is mentally constructed. This holds for all three aspects of time: order, duration, and temporal perspective (i.e. past, present and future). It follows that time as represented by the real numbers linearly ordered by the 'earlier than' relation, together with the (presumably non-physical) concept 'now', is not likely to provide a faithful representation of tense and aspect. Instead, it is proposed that the future is cognitively represented as a set of goals to be attained, together with possible courses of action to achieve those goals. It has been claimed Suddendorf and Corballis [1997] that even remembering the past (i.e episodic memory) is a by-product of a more general capacity to imagine possible worlds, which finds its main use in contemplating alternative courses of action.

> The plan unites the past (a desired state) with the present (an attempt) and the future (the attainment of that state) . . . causality and planning provide the medium through which the past is glued to the present and the future. (Trabasso and Stein Trabasso and Stein [1994])

The view that tense and aspect relate to planning is not new; it was expressed by Steedman in [Steedman, 1997, p. 932–3] as

> The semantics of tense and aspect is profoundly shaped by concerns with goals, actions and consequences . . . temporality in the narrow sense of the term is merely one facet of this system among many. Such concerns seem to be the force that determines the logic that is required to capture its semantics as the particular kind of dynamic system outlined above, whose structure is intimately related to knowledge of action, the structure of episodic memory, and the computational process of inference.

Steedman similarly believed that temporal reasoning formalisms from AI would be very useful in this context, but the formalisms then available were somewhat cumbersome. The event calculus as reformulated in constraint logic programming[8] provides a much more flexible tool.

The next section illustrates the role of goals and planning in the English future tense; but before we come to this we have to emphasize the important role of *perspective* when talking about eventualities. These are not intended to be ‘things out there’, but ways of viewing and structuring the world. Thus, if one defines following Comrie, imperfective aspect as involving the ‘internal temporal contour of a situation’, this should not be read realistically, but as a particular construction of an event. The same chunk of space-time can be viewed perfectively, i.e. without internal structure, and imperfectively, for instance with a structure of goal, consequent state, and actions leading up to the goal.

---

[8] We will not dwell on the technical differences between standard logic programming (as enshrined in Prolog), and constraint logic programming. The latter is more general; for instance it allows the unification of two constants.

3.3. The importance of being goal-oriented: future tense

The English future tenses make the connection between tense and the structure of plans particularly clear, because they can be seen to codify various ways of achieving a goal.

(5) The sun rises at 6:30am tomorrow.
(6) a. Bill will throw himself off the cliff.
  b. Bill is going to throw himself off the cliff.
(7) a. I will fly to Chicago tomorrow.
  b. I am going to fly to Chicago tomorrow.
  c. I was going to fly to Chicago tomorrow, but my boss forbade me.
(8) a. *I go to Chicago unless my boss forbids me.
  b. (Google) I am going unless some unknown demand stops me.
  c. (Google) I will go unless there is severe or dangerous weather.
(9) a. *I fly to Chicago if my boss asks me.
  b. ?*I am going if you go.
  c. I am going if my health allows me/if I am able.
  d. (Google) Barak said to Deborah, "I will go if you go with me. I will not go if you don't go with me."

Syntactically, future tense can thus be expressed by simple present (5) (although not in all contexts, cf. (8-a), (9-a)), futurate progressive (8-b), (9-c) (again not in all contexts, cf. (9-b)), and with the help of the auxiliaries *will* and *be going to*, which have fewer restrictions than the aforementioned constructions: compare for example (9-a) and (9-b) with (9-d).

Semantically, one can distinguish two main dimensions along which future events can be classified. The first dimension concerns two possible perspectives on future events in so far as they can be affected by humans: as events per se, and as goals, to be achieved by a plan (which may possibly fail). In very rough outline one may say that the use of the present tense emphasizes the first perspective. A good example of this is (5).

Examples (8-a) and (9-a) show that the present tense is no longer allowed if even a mild form of conditional planning is introduced. By contrast, sentences (8-c) and (9-d) show that the auxiliary *will* is fine with planning. Indeed, the auxiliaries often indicate that some amount of planning is involved, but here an orthogonal dimension comes into play.

Suppose we view a future event from the perspective of goals and plans. If *will* is used in contexts such as (7-a), it is indicated that no actions of self interfere with the execution of the plan. On the other hand, if *be going to V* is used in that same context ((7-b) and (7-c)), the possibility of an obstacle arising is deliberately left open. Thus sentence (6) is false if Bill in the end does not jump off the cliff, unlike sentence (6-b): as Comrie [Comrie, 1976, p. 64–5] remarks, the second sentence can be shouted as a warning and an injunction to do something to prevent Bill from jumping, whereas the first sentence cannot be used in this way. This does not mean that if *will* is used, no actions or events can interfere with the achievement of the goal; rather, these actions and events must be mentioned explicitly in subordinate clauses, as in (8-c) or (9-d). In the case of *will*, no obstacles are envisaged apart from those explicitly listed, whereas *be going to* can be used much more freely. A good instance of this is (10)

(10) (Google) Tony Blair in 1997: 'I am going to be a lot more radical in government than people think'.

3.4. A glimpse of the formalism

The preceding considerations show that a formal semantics for tense and aspect may take the form of a planning formalism which is able to talk about goals and actions, and which includes a theory of causality together with a principle of inertia. Such a formalism is presented in van Lambalgen and Hamm [2004]. It consists of an 'event calculus' which has found applications in robotics, here reformulated using the computational machinery of constraint logic programming.

The reader may well wonder why robotics can provide a source of inspiration to linguistics. The reason can be found in the nature of robotic computation.[9] A typical computation in robotics proceeds as follows. A goal is specified, which can be a certain location (say in an office building) and an action to be performed at that location (e.g. pick up outgoing

---

[9] If one also assumes that human path planning shares the main computational features with robot path planning, one may speculate about the origin of language in motor programs, as some indeed have done (e.g. Arbib, Corballis, Greenfield).

mail). Next a plan is computed, that is, a sequence of actions to get the robot to the required location, which can be obtained by backward chaining from the goal to obtain a sequence of subgoals, the last one of which can be executed in the robot's initial position and state. Such a computation requires a world model (including a map of the building, a causal theory of the robot's actions, a specification of values of variables such as 'door open/closed', the initial position and state of the robot, a record of its past and current actions, . . .), a repertoire of activities and actions (e.g. 'follow wall', 'go through door') and of possible observations (e.g. 'door open/closed'). On the basis of the world model a plan is computed. While the robot executes the plan, it registers its observations of the world and its actions in the world model; knowledge of its actions may be important for the robot to estimate its current position. The plan may have to be recomputed in mid-course when the world model must be updated due to new observations (e.g. of a closed door which was expected to be open on the basis of the initial world model, or a wrong estimate of the current position). Note that a plan may consist of continuous activities ('traverse distance x at speed y') and (almost) instantaneous actions ('get sonar reading'), so that the latter take place during the former.

This description should be sufficiently suggestive to enable the reader to see the connection with linguistic processing. The listener starts with an initial discourse model, in which a newly arriving sentence must be integrated computably. Suppose the main verb of the sentence is nonstative. If the sentence is in one of the simple tenses, it is unpacked in an action and its participants, and the discourse model is updated accordingly. This is the analogue of updating the world model with representations of individuals and actions. In more complex cases, such as (6-a) and (6-b) above, the sentence expresses the existence of a plan directed toward the goal formulated in the VP. If on the contrary the main verb of the sentence is stative, the sentence can be viewed as analogous to an observation report, and the discourse model is accordingly updated with a property.

Since the formalism is unfortunately much too involved to explain in full formal detail in the space allotted to us, we first give a qualitative description, which will then be illustrated in 3.5 by equally qualitative sketches of the computations of the formalism as applied to the examples in section 3.1.

3.4.1. The language of the event calculus

The event calculus is a planning formalism which allows one to talk about actions, goals and causal relations in the world. Its main function is to return a plan given a goal, the initial state, and causal relationships. Formally, the event calculus is a many-sorted logic. It has two different sorts for events viewed either perfectively or imperfectively.[10] The former are called event types and are symbolized by $e, e', \ldots, e_0, \ldots$. The latter are called fluents[11], and symbolized by $f, f', \ldots, f_0, \ldots$. One may think of the event types as action types, such as for example 'break' or 'ignite'; the fluents can be thought of as time-varying properties, for example 'being broken' or 'walking'; the time-parameter in fluents is implicit, but they can have further parameters (e.g. for the subject of 'walking'). The real distinction between event types and fluents comes from the different roles they play in the axioms of the event calculus.

Continuing with the ontology, we note that the universe must also contain sorts for individuals ('John'), for real numbers interpreted as instants of time, and for various other real quantities (e.g. position, velocity, degree of some quality). The word 'ontology', while referring only to the domains of models, is potentially misleading here. Aspect is not concerned with the real temporal constitution of events, whatever that may mean, but with our construal of events. Use of the perfective aspect does not mean that an event is inherently completed, only that we view it as completed. Likewise, the different kinds of eventualities introduced in the event calculus are not conceived of as being 'out there' in the world, but just different ways in which we conceptualize the world.

The primitive predicates may look somewhat baroque, but they comprise the bare minimum necessary to talk about two forms of causality, instantaneous (as in two balls colliding) and continuous (as when a force is acting). Here we list only the predicates for instantaneous change:

---

[10] Compare this surprising remark from Jean-Yves Girard, in a paper blasting classical logic: 'Il y a d'autres intuitions de base qui ont été évacuées par la logique, ainsi la distinction essentielle entre *parfait* et *imparfait*, distinction rendu en français par le choix des temps, en russe par le changement de verbe. Cette nuance n'existe pas dans le monde vériste.' (*La logique comme géométrie du cognitif*)

[11] Newton's name for variables depending on time.

– *Initially*(f) ('fluent f holds at the beginning of the discourse')

– *Happens*(e, t) ('event type e has a token at t')

– *Initiates*(e, f, t) ('the causal effect of event type e at time t is the fluent f')

– *Terminates*(e, f, t) ('the causal effect of event type e at time t is the negation of the fluent f')

– *Clipped*(s, f, t) (roughly, 'an event type terminating f has a token between times s and t')

– *HoldsAt*(f, t) (fluent f is true at t; see also below)

All semantically relevant causal relations have to be translated into this language. Thus, if we consider the French sentences

(11) Jean appuya sur l'interrupteur. La lumière l'éblouissait.
Jean pushed the button. The light blinded him.

we notice an event type 'appuyer sur l'interrupteur' which has a causal effect on the state of the light. In order to proceed further, one needs an automatic procedure to translate natural language expressions (e.g. 'x appuyer sur l'interrupteur' into the formal language of the event calculus (a term $e(x)$). In broadest outline, this goes as follows. Event types and fluents are terms which can be seen as codes for formulas via *reification* (also called Gödelization). First represent a verb by a predicate $A(\bar{x}, t)$ (all free variables exhibited; t is the temporal parameter). One may form the expression[12] $\{t \mid A(\bar{x}, t)\}$ (with the $\bar{x}$ as free parameters) – one may think of this expression as the fluent $f(\bar{x})$, which thus contains an implicit temporal parameter. In order to enforce the interpretation of the fluent $f(\bar{x}) = \{t \mid A(\bar{x}, t)\}$ as a set, the event calculus uses the truth predicate[13] $HoldsAt(f(\bar{x}), s)$, intuitively meaning $s \in \{t \mid A(\bar{x}, t)\}$, i.e. $A(\bar{x}, s)$.

Event types $e(\bar{x})$ can be constructed as (the Gödel number of) a formula $\exists t A(\bar{x}, t)$, i.e. abstracting away from time.[14]

---

[12] This assumes that the language from which A is taken has a pairing operation; Chapter 6 of van Lambalgen and Hamm [2004] has the details.

[13] To ensure that *HoldsAt* really is a truth predicate, axioms for *HoldsAt* must be added to those of the event calculus. These are nontrivial (but usually omitted in the literature); again see Chapter 6 of van Lambalgen and Hamm [2004].

[14] This construction explains why event types and fluents are also suitable for representing perfect and imperfect nominals, respectively, since from the former vestiges of time have been eradicated, but not the latter. See Chapter 12 of van Lambalgen and Hamm [2004].

3.4.2. The computational machinery of the event calculus

One important difference between (neo-)Davidsonian event semantics and the approach presented here is that we relate the chosen primitives to each other by means of axioms, whereas the former approach to semantics relies on an intuitive understanding of its predicates *Cul*, *Holds*,... One reason for this is the general methodological principle that one can derive exact predictions only from formalized theories, also in semantics. However, by far the most important reason is that we want to explain how during language comprehension an event structure, more generally a discourse model, is computed. We note again that not only Davidsonian event semantics, but also work in the tradition of Cognitive Linguistics stands in need of such a computational approach to substantiate its claims of cognitive relevance. We will indicate here, albeit informally, how axioms for events, together with a suitable nonmonotonic logic can achieve this.

Here is an example of an axiom:

> if a fluent f holds *initially* or has been *initiated* by some event occurring at time t and no event *terminating* f has occurred between t and $t' > t$, then f *holds at* t'

This axiom expresses one form of temporal inertia, analogous to Newton's first law: if the fluent f starts to hold at t then it will continue to hold uninterruptedly from t to t' unless an *explicit* cause terminates f in the meantime. To see what this axiom contributes to the construction of the event structure, consider our simple French example

(12) Il faisait chaud. Jean ôta sa veste. (Imp, PS)
It was hot. Jean took off his sweater.

Intuitively, this narrative determines an event structure in which 'hot' acts as a background which is true all the time; the foregrounded event ('taking off one's sweater') is placed inside this background. One arrives at this structure by means of the following argument. World knowledge contains no causal link to the effect that taking off one's sweater changes the temperature. Since it *is* hot at some t before *now*, the state *hot* must either hold initially or have been initiated. The latter requires an event, which is however not given by the discourse. Therefore *hot* holds initially. Similarly no terminating event is mentioned, so that *hot* extends indefinitely, and it follows that the event described by the second sentence must be positioned inside *hot*.

However, in the above explanation of the effect of the axiom we have relied on the tacit understanding that, since the discourse itself mentions only two eventualities, one perfective and one imperfective, the event structure determined by the discourse contains these two events only. But neither the discourse nor the axioms enforce that this is so: the addition of further events (e.g. corresponding to ‘et Marie ouvrit la fenêtre’) does not contradict the axioms, while possibly changing the event structure. Indeed, if we also add an atomic sentence to the effect that opening the window terminates it being hot inside, the event structure becomes different. Therefore the axiom has the desired effect only in ‘minimal’ models of the discourse, where ‘minimal’ here refers to the dual requirement that

i. the model only contains those occurrences of events forced to be there by the discourse and the axioms
ii. the interpretation of the primitive predicates (*Initiates* etc.) is as small as is consistent with the discourse and the axioms

Speaking informally still, this means that in a model of the discourse, no unforeseen events are allowed to happen, and similarly that all causal influences are as expected. These are of course the same requirements that we found to be important while discussing planning in section 3.2. There is no need to explain the concept of ‘minimal model’ here beyond the intuitive level, but we should note that the choice to work with a minimal model instead of all models leads to *nonmonotonicity* in the construction of discourse models. That is, extending a discourse with a new sentence (as in ‘et Marie ouvrit la fenêtre’) may invalidate a conclusion derived from the original discourse, in this case that the state of it being hot inside extends indefinitely beyond the event time determined by the sentence ‘Jean ôta sa veste’; given that opening the window has a cooling effect, the state of it being hot will be terminated. In Chapters 9 and 11 of van Lambalgen and Hamm [2004] it is argued that it is precisely the possibility to retract previously inferred conclusions which allows a rigorous treatment of the semantics of the English progressive and of coercion. Below, in section 3.8, we will discuss the implications of this proposal, with its recurrent recomputations of discourse models, for the interpretation of brain signals.

The most important metatheorem about the formalism is that minimal models exist[15], *and can be computed efficiently*. This is a consequence of the syntactic structure of the axioms of the event calculus and the formulas used to translate sentences in a discourse, which allows one to use the techniques of constraint logic programming, in particular its version of the Herbrand models of ordinary PROLOG.

3.5. Computing event structures for (PS, Imp) combinations

We will now further illustrate the formalism's operation by tracing the computations involved in determining the event structures for the remaining two French examples. It will be seen that in these cases, the principle of inertia, as embodied in the axioms of the event calculus, together with the minimization procedure described above, jointly produce the required event structure.

(13) Jean attrapa une contravention. Il roulait trop vite. (PS, Imp)
Jean got a ticket. He was driving too fast.

This example dates from the bygone days when speeding cars were stopped by the police instead of being photographed. It is given that the event of getting a ticket occurred sometime in the past. It is also given that the fluent *speeding* was true some time in the past, hence it holds initially or has been initiated. We have to determine the relative position of event and fluent. World knowledge yields that getting a *ticket* terminates, but not initiates, *speeding*. Since this is the only event mentioned, *speeding* holds from the beginning of discourse, and is not re-initiated once it has been terminated.

In the second example (14) the same order of the tenses yields a different event order, guided by the application of causal knowledge.

(14) Jean appuya sur l'interrupteur. La lumière l'éblouissait. (PS, Imp)
Jean pushed the button. The light blinded him.

---

[15] There may exist other models as well (also ones which are not computable), but these are taken to be irrelevant for the representation of discourse. Minimal models can be obtained as the least fixed point of a suitable consequence operator.

One (occurrence of an) action is mentioned, *push*ing the light button, which has the causal effect of initiating the light being on when its current state is off. No terminating event is mentioned, so that the light remains on. It also follows that the light must be off for some time prior to being switched on, and therefore that it must be off at the beginning of discourse. The definite article in '*La* lumière' leads to a search for an antecedently introduced light, which successfully terminates after unification with the light introduced in the first sentence; therefore it is this light which is too bright.

3.6. Computing event structures more formally: integrity constraints

To conclude this section, we return to one of our starting points, introduced in section 3.1: *we view a sentence* S *as a goal ('make* S *true') to be achieved by updating the discourse model*. It is our purpose here to make this notion of update formally precise, in preparation for the application to DRT in the next section. Consider the following example from van Lambalgen and Hamm [2004] involving the English perfect.

(15) I have caught the flu.

Let f be the fluent expressing *having the flu* and let e be the corresponding infection event. Assume that the event and the fluent f are related by the following formula of the event calculus: $\mathit{Initiates}(\mathrm{e}, \mathrm{f}, \mathrm{t})$, expressing that world knowledge that the event e is a cause of the fluent f. Informally, sentence (15) is true if I have the flu *now*. This just states the often observed fact that the English perfect has present relevance. Now let us see what is involved in viewing (15) as the goal 'Make "I have caught the flu" true in the given discourse model'.

Assume we are given a discourse model, say presented in the form of a list of facts concerning events and fluents. We have to construct a (minimal) adaptation of the discourse model in which $\mathit{HoldsAt}(\mathrm{f}, \mathit{now})$ is true. This is not just a matter of adding $\mathit{HoldsAt}(\mathrm{f}, \mathit{now})$, since the truth of this sentence in the model might have further consequences for the model, as it has in the case at hand. The sentence $\mathit{HoldsAt}(\mathrm{f}, \mathit{now})$ is therefore taken to trigger a kind of abductive reasoning using axioms of the event calculus and if necessary knowledge about the world. Applied to our case, this reasoning proceeds as follows. Remember that in section 3.4 the following axiom of the event calculus was introduced:

If a fluent f holds *initially* or has been *initiated* by some event occurring at time *t* and no event *terminating* f has occurred between t and $t' > t$, then f *holds at* time t'.

We know that the fluent f expressing that I have the flu is initiated by the event e. No terminating event for f has been mentioned, so we conclude by a version of closed world reasoning that no such event occurred. According to the above axiom there is only one fact missing to establish $\mathit{HoldsAt}(\mathrm{f}, \mathit{now})$, namely that the infection event e actually occurred before *now*: $\mathit{Happens}(\mathrm{e}, \mathrm{t})$, $\mathrm{t} < \mathit{now}$. We therefore add this fact and its logical consequences (and nothing else) to the model. The resulting model is then a model in which $\mathit{HoldsAt}(\mathrm{f}, \mathit{now})$ is true. By this kind of abductive reasoning we therefore get the inference from (15) that an infection event occurred in the past.

Formally, the above reasoning process is carried out by the derivation procedure used in logic programming called *resolution*. A derivation is started with some formula as the top query; in our case this formula is $\mathit{HoldsAt}(\mathrm{f}, \mathit{now})$. To emphasize that the purpose of the derivation is to see whether the formula can be realized in a model, it is written as $?\mathit{HoldsAt}(\mathrm{f}, \mathit{now})$. The resolution process involves matching the consequent of a suitable axiom to the given clause, and replacing the top query by the antecedent of the axiom. The process is repeated as long as there are consequents of axioms to be matched. This process will usually end with a query that cannot be further resolved; in the case above with $\mathit{Happens}(\mathrm{e}, \mathrm{t})$, $\mathrm{t} < \mathit{now}$. In ordinary resolution the top query would now be considered as failed. It is however also possible to interpret the final query as an instruction to update the model with the ingredients necessary to satisfy the top query. Read in this way, the top query is called an *integrity constraint*, and is written as

$$?\mathit{HoldsAt}(\mathrm{f}, \mathit{now})\ \texttt{succeeds}.$$

The upshot of this discussion is therefore that the update-character of tenses is represented formally by means of integrity constraints.

A second type of integrity constraint that is useful in this context is one in which the top query must fail. Consider the following example of the simple past tense

(16) John ran.

Let e be the 'run' event type (constructed formally as indicated in section 3.4.1), then a first stab at a formalization would be

$$?\mathit{Happens}(\mathrm{e},\mathrm{t}),\ \mathrm{t} < \mathit{now}\ \texttt{succeeds}.$$

This formulation still allows the possibility that running is going on now, which is undesirable for activities. Here it is important that e is an event *type*, which may have different tokens in any given model. In the minimal models which are of interest to us, the set $\{\mathrm{t} \mid \mathit{Happens}(\mathrm{e},\mathrm{t})\}$ is a finite set of intervals; each of these intervals constitutes a token of e. One does not want to exclude that one of these tokens lies in the future, but there must be one token which is completely in the past.[16] This can be achieved by means of the negative integrity constraint

$$?\mathit{Happens}(\mathrm{e},\mathit{now})\ \texttt{fails}.$$

This says that any update resulting in $\mathit{Happens}(\mathrm{e},\mathit{now})$ is forbidden. The simple past is thus represented by means of two integrity constraints, one positive (demanding success), and one negative (demanding failure). From now on we will use the following terminological convention: an integrity constraint IC is said to be *satisfiable* if it can be made to succeed in case it is positive, and can be made to fail in case it is negative. It is also possible to combine the two integrity constraints into one by writing

$$?\mathit{Happens}(\mathrm{e},\mathrm{t}),\ \mathrm{t} < \mathit{now},\ \mathit{Happens}(\mathrm{e},\mathit{now})\ \texttt{succeeds}.$$

One last point about the simple past. It has often been observed that the past tense is anaphoric in the sense that it needs to be anchored in an antecedently given context. Thus, by itself 'John ran' is not felicitous; it becomes felicitous when an adverbial like 'On Wednesday, ...' is added. In the present framework, this feature can be captured by adding a fluent for

---

[16] The reader might object that a token of run could occur now. Note however that in such cases one must use the present progressive, which must be represented by a fluent, not an event. (See Chapter 10 of van Lambalgen and Hamm [2004].)

the context. Let f be a new fluent-constant, not yet present in the discourse, then the past tense is represented by two integrity constraints of the form

$$?\mathit{HoldsAt}(\mathrm{f},\mathrm{t}),\ \mathit{Happens}(\mathrm{e},\mathrm{t}),\ \mathrm{t} < \mathit{now}\ \texttt{succeeds},$$

and

$$?\mathit{HoldsAt}(\mathrm{f},\mathrm{t}),\ \mathit{Happens}(\mathrm{e},\mathrm{t}),\ \mathrm{t} \geq \mathit{now}\ \texttt{fails}.$$

The constant f is then to be unified with material from the context. (This procedure will be illustrated in section 4 on DRT.) If necessary, further conditions on f and e can be expressed using integrity constraints, for instance two conditions used in DRT's representation of the Simple Past, namely that the context lies entirely in the past[17], and that the event lies entirely within the context. The first condition can be expressed by the integrity constraint

$$?\mathit{HoldsAt}(\mathrm{f},\mathrm{t}),\ \mathrm{t} \geq \mathit{now}\ \texttt{fails},$$

and the second by

$$?\neg\mathit{HoldsAt}(\mathrm{f},\mathrm{t}),\ \mathit{Happens}(\mathrm{e},\mathrm{t})\ \texttt{fails}.$$

3.7. *Consequences for the theory of meaning: compositionality*

It should be said here, although we lack the space to elaborate on the topic, that the emphasis on nonmonotonic reasoning processes in the construction of meanings leads to a theory of meaning which is very different from the standard picture. In the latter, it is assumed that there are atomic units of meaning, i.e. expressions whose meaning is independent of the context in which they occur, and which are combined as syntax dictates to form meanings of compound expressions. It is often argued that both production and comprehension of a potentially infinite set of sentences needs this form of compositionality. We disagree here. Indeed, the procedure for the computation of discourse models or event structures outlined above is very much top-down, and most expressions can change

---

[17] This condition is somewhat doubtful though, as witnessed by Google's 45 million hits for the expression 'Today, I went to ...'.

their meaning as the discourse context in which they occur dictates. The reader may consult Chapter 11 of van Lambalgen and Hamm [2004] for more on this topic in connection with the phenomenon of coercion.

3.8. Consequences for brain imaging of language processing

The ideas presented above lead to rather straightforward predictions concerning semantic processing, predictions which can be tested using electrophysiological methods such as ERP.[18] The main ideas can be illustrated using the imperfective paradox, the observation that, while activities in the progressive tense are generally veridical, accomplishments are not. The following are examples of sentences containing progressivized activities (ACT) and accomplishments (ACC):

(17) a. John was running in the campus when he saw his friend Paul. (ACT+)
b. John was running in the campus when he was hit by a car. (ACT−)
c. John was crossing the street when he saw his friend Paul. (ACC+)
d. John was crossing the street when he was hit by a car. (ACC−)

Examples (17-a) and (17-c) entail that John ran in the campus and that John crossed the street respectively. The event described by the *when* clause does not terminate either the activity of running or that involved in crossing the street, that is walking from one side of the street to the other. As a result, it does not prevent the goal of reaching the other side of street in (17-c) from being attained. An asymmetry between activities and accomplishments is introduced by a manipulation of the event described by the *when* clause, as exemplified by (17-b) and (17-d) above. While (17-b) entails that John ran in the campus, although the fact that he was hit by a car presumably terminated the running activity, (17-d)

---

[18] This section describes experiments which are under way at the F.C. Donders Centre for Neuroimaging (Nijmegen, The Netherlands). We are greatly indebted to Giosue' Baggio (who is designing these experiments) for help with this section.

entails that John did not reach the other side of street. The termination of the activity of walking from one side of street to the other implies that the goal state (i.e. having reached the other side) was not attained. The interpretation of (17-d) can be seen as an instance of nonmonotonic reasoning. Initially a minimal model of the progressivized clause is computed, entailing that the event described the corresponding perfective clause occurred. Augmenting the model with new information, such as that provided by the *when* clause in (17-d), destroys the former inference. In this project we intend to investigate using EEG/ERP the effects of non-monotonicity on the interpretation of sentences containing activities and accomplishments in the progressive tense.

The theory of semantic processing presented in van Lambalgen and Hamm [2004] and informally explained above predicts that the model computed during the first clause of ACC− sentences like (17-d) is readjusted to accommodate the inference that the goal state (i.e. reaching the other side of the street) was not attained. Compared to the three other conditions, the recomputation of discourse models is expected to elicit a larger anterior negative deflection, especially during the second clause. Power increases in the gamma band have been related to the integration of world knowledge into a discourse representation Hagoort et al. [2004]. The interpretation of accomplishments in which the goal state is prevented from being attained involves knowledge of the causal relations between the events described. Therefore, we expect that ACC− sentences induce a significant activity increase in the gamma frequency range, possibly with an anterior scalp distribution. As to the neural sources of these effects, the left inferior prefrontal cortex (LIPC), and in particular Brodmann's areas 45 and 47 Hagoort et al. [2004], might be crucial for the interpretation of ACC− sentences like (17-d). A further question, again to be addressed using EEG source analysis or fMRI, is whether the LIPC is the only area recruited by the recomputation of discourse models (i.e. by nonmonotonic reasoning) or whether other anterior brain regions are involved. A prediction following from the theory proposed in van Lambalgen and Hamm [2004] is that planning areas such as the frontal lobes Koechlin et al. [1999] have a critical role in readjusting discourse models and are therefore implied in the interpretation of accomplishments. A theoretical model based on the picture just sketched predicts a significant interaction of the factors aspectual class and event type, reflecting the distinctive EEG signal elicited by the recomputation of models in ACC– sentences.

4. DRT and event calculus

It is the purpose of this section to argue that DRT and semantics based on the event calculus can be of mutual benefit, and that the event calculus allows one draw out some of the cognitive implications of DRT in more explicit form. On the one hand, techniques have been developed in DRT for transforming syntactic representations into semantic ones[19]. In particular DRT furnishes a systematic device for disambiguating natural language expressions. In van Lambalgen and Hamm [2004], semantic representations are constructed based on cognitive considerations involving the mental representation of time and action. This leaves open the question, though, how semantic representations are determined on the basis of the linguistic input. If one assumes that the processing of an utterance starts by analyzing it syntactically, one needs a mechanism to hook up the syntactic analysis to the integrity constraints[20]. We will show in section 4.1 by means of several examples that one can think of the integrity constraints as being derived from DRSs; put in different terms, DRSs can be embedded in the minimal models derived from integrity constraints. In a way this provides an existence proof for an algorithm from syntax to integrity constraints. More generally, it shows that the two approaches are compatible. In fact, their compatibility can take two different forms. On the one hand we can use DRSs which represent discourses (and whose construction has already benefitted from the principles of dynamic interpretation that are central to DRT) as input to an algorithm which converts them into constraint logic programs and then use these clp's as premises for (nonmonotonic) inferencing within constraint logic programming, taking advantage of its special computational properties. But at least in some instances it is also possible to convert DRSs for

---

[19] From this point on we assume the reader is familiar with the rudiments of DRT.

[20] There are alternatives to this assumption, such as the usage-based approaches of Goldberg [1995], Tomasello [2003] and Verhagen [2005]. The problem of relating surface structure to semantic representation takes on a different form here.

the individual sentences which make up a discourse severally into clp's and to compute the intersentential connections which give the discourse its cohesion (such as trans-sentential temporal and nominal anaphora, etc.) at the level of the logic programs and with the methods it provides.[21]

In this final section of the paper we can do no more than give the merest indication of what these two options come to. To determine exactly how far either option will go is a non-trivial matter. The difficulty is directly connected with a fundamental tension that exists between our computational approach and DRT, the important similarities in general outlook, to which we have drawn attention in the preceding sections, notwithstanding. The representation formalisms used in DRT (the so-called 'DRS-languages') are motivated by the desire to represent, in a logically transparent way, whatever information is expressible in natural languages and to do so in forms that remain as close as possible to those in which that information is expressed in the natural language in question. DRS-languages are logically transparent in the sense that they come with a well-defined model theory, including a strictly compositional truth definition. The price these formalisms pay for their expressive power is that they are in general not 'computable': there are at best partial algorithms for deriving logical consequences from given 'premise' representations. Moreover, no systematic theory of nonmonotonic reasoning from such representations has so far been developed, and we doubt that this could be done in any natural and direct way. One important claim of this paper is indeed that the event calculus provides the ideal representational format for this purpose. There are two main reasons for this: (1) because DRSs can be translated into integrity constraints, the inference mechanisms of logic programming can be used to generate an inference mechanism for DRSs, and (2) the axioms and inference mechanisms of the event calculus are concerned with change in time, and therefore allow one to compute the development of a DRS over time. If John visits Mary at 1:00pm, and is told by her colleague that she is out for lunch, how is John going to use that information at 2:00pm? Linguistic information

---

[21] In the cases where it is formally possible to follow either of these routes, the question arises which of them is the cognitively more realistic one, but this is a question to which it would be premature for us to venture an answer.

must always be integrated with world knowledge, in particular with how objects and events behave over time.

It is precisely these computational concerns which have been the central motive behind the choice of constraint logic programming and the event calculus as a formalism for the construction and further manipulation of semantic representations, along the lines of van Lambalgen and Hamm [2004]. But the desirable computational properties of the formalism depend crucially on the special restrictions imposed on its syntax. This means in particular that there is no simple algorithm for turning arbitrary DRSs into logic programs and integrity constraints. The conversion is straightforward for simple DRSs (i.e. DRSs which do not contain complex DRS conditions, see Kamp and Reyle [1993]). For complex DRSs the conversion is not straightforward, and at the present it is not clear to us exactly when conversion is possible without loss of content.

4.1. From DRSs to integrity constraints

We will show by means of three examples of increasing complexity how integrity constraints can be read off from DRSs. Since in this section we combine two formal systems – DRT and event calculus – we will use slightly different notations for terms in the two subsystems. For instance we will write *e* for an event term in DRT and e for event terms in the event calculus framework.

4.1.1. Single DRS

Let us start with a very simple example.

(18) Max arrived.

The DRS for this sentence is:

(19) $$\boxed{\begin{array}{ccc} m & t & e \\ Max(m) & t < n & e \subseteq t \\ \multicolumn{3}{c}{e : arrive(m)} \end{array}}$$

Here m, t and e are the discourse referents of the above DRS; t is the contextually determined reference time (supposed to lie in the past), and e is the event time.

Since DRSs introduce existential presuppositions which have to be accommodated, integrity constraints as introduced in section 3.6 are the appropriate means to translate DRSs into logic programming. We will now show how to write an integrity constraint for DRS (19).

For this we assume that predicates $\mathit{max}(\mathrm{x}, \mathrm{t})$ and $\mathit{arrive}(\mathrm{x}, \mathrm{t})$ are given. These predicates will be used in their reified forms which are derived via the procedure given in section 3.4.1. By transforming predicates into terms we get expressions which can be used as arguments of the predicates of the event calculus. The first possibility to derive terms via reification, applied to $\mathit{max}(\mathrm{x}, \mathrm{t})$, results in the fluent term $\mathit{max}[\mathrm{x}, \hat{\mathrm{s}}]$ which can be used as argument of the *HoldsAt*–predicate. The second possibility for reification, applied to $\mathit{arrive}(\mathrm{x}, \mathrm{t})$, derives the event type $\exists \mathrm{s}.\mathit{arrive}[\mathrm{x}, \mathrm{s}]$, a term which can be used as argument of *Happens*. We now show how to represent the existential presupposition of the DRS, the temporal contribution of the past tense, and its anaphoric character, in a single integrity constraint. Let f be an unspecified context fluent[22] anchoring the reference time; this is formalized as a clause $\mathit{HoldsAt}(\mathrm{f}, \mathrm{t})$. The discourse referent e corresponds to event type $\exists \mathrm{t}.\mathit{arrive}[\mathrm{x}, \mathrm{t}]$, n to *now* and t to the context fluent f. We therefore get as a first translation of the DRS the following integrity constraint

(20) $?\mathit{HoldsAt}(\mathrm{f}, \mathrm{t}),\ \mathit{HoldsAt}(\mathit{max}[\mathrm{x}, \hat{\mathrm{s}}], \mathrm{t}),\ \mathit{Happens}(\exists \mathrm{s}.\mathit{arrive}[\mathrm{x}, \mathrm{s}], \mathrm{t}),$
$\mathrm{t} < \mathit{now},\ \neg \mathit{Happens}(\exists \mathrm{s}.\mathit{arrive}[\mathrm{x}, \mathrm{s}], \mathit{now})$ `succeeds`

At the end of section 3.6 we indicated how to capture the (possibly overly strong) condition $t < n$ by means of the integrity constraint $?\mathit{HoldsAt}(\mathrm{f}, \mathrm{t}),\ \mathrm{t} \geq \mathit{now}$ fails, and analogously for the condition $e \subseteq t$. This completes the translation of this simple DRS into the language and inference mechanisms of the event calculus. From now on we will skip as much as possible from the internal structure of event types and fluents and will present integrity constraints as informally as possible.

---

[22] The context fluent is formally represented by a new constant, which may then be identified with fluents available in the discourse.

4.1.2. Merging DRSs

The following example (taken from von Genabith et al. [2005]) will allow us to indicate the computational treatment of anaphora in this framework.

(21) A delegate arrived. She registered.

The DRS for the first sentence is given in (22).

$$(22)\quad \boxed{\begin{array}{ccc} x & t & e \\ \mathit{delegate}(x) & t < n & e \subseteq t \\ & e : \mathit{arrive}(x) & \end{array}}$$

Write h for the context fluent, f(x) for the predicate ‘delegate(x)’, e(x) for the event ‘arrive(x)’, e$'$(x) for the event ‘register(x)’. The integrity constraint corresponding to (22) is then given by

$$(23)\quad ?\mathit{HoldsAt}(\mathrm{h}, \mathrm{t}),\ \mathit{HoldsAt}(\mathrm{f}(\mathrm{x}), \mathrm{t}),\ \mathit{Happens}(\mathrm{e}(\mathrm{x}), \mathrm{t}),\ t < \mathit{now}, \\ \neg\mathit{Happens}(\mathrm{e}(\mathrm{x}), \mathit{now})\ \texttt{succeeds}$$

possibly together with a negative integrity constraint to express that the context is in the past. Since we intend to concentrate on the formalization of anaphoric resolution, we will not write the full set of integrity constraints. The DRS for the second sentence in (21) is:

$$(24)\quad \boxed{\begin{array}{ccc} y & t & e \\ & t < n & e \subseteq t \\ & e : \mathit{register}(y) & \end{array}}$$

We now need a further context fluent h$'$. The most interesting feature of this sentence is the need to represent anaphoric ‘she’. In line with DRT’s representation of individuals as predicates, we opt to represent ‘she’ as a new fluent variable s(x), to be unified with given material. The integrity constraint corresponding to (24) is then given by

$$(25)\quad ?\mathit{HoldsAt}(\mathrm{h}', \mathrm{t}),\ \mathit{HoldsAt}(\mathrm{s}(\mathrm{x}), \mathrm{t}),\ \mathit{Happens}(\mathrm{e}'(\mathrm{x}), \mathrm{t}),\ t < \mathit{now}, \\ \neg\mathit{Happens}(\mathrm{e}'(\mathrm{x}), \mathit{now})\ \texttt{succeeds}$$

DRT allows one to merge the two DRSs (22) and (24) into a single DRS expressing the information contained in both. This DRS can be translated into an integrity constraint as above. We will show here that the integrity constraint derived in this manner is the same as the one obtained by fusing (23) and (25) and and applying anaphoric resolution[23]. Clearly the context h' for (25) is furnished by the query in (23). We have to ask the reader to take on trust that this can be represented formally by writing a program clause defining h'

(26) $$HoldsAt(\mathrm{h},\mathrm{t}) \wedge HoldsAt(\mathrm{f}(\mathrm{x}),\mathrm{t}) \wedge Happens(\mathrm{e}(\mathrm{x}),\mathrm{t}) \rightarrow HoldsAt(\mathrm{h}',\mathrm{t})$$

The clause (26) can be used to reduce the query in (25) via resolution, which yields the new integrity constraint

(27) $$?HoldsAt(\mathrm{h},\mathrm{t}),\ HoldsAt(\mathrm{f}(\mathrm{y}),\mathrm{t}),\ Happens(\mathrm{e}(\mathrm{y}),\mathrm{t}),\ HoldsAt(\mathrm{s}(\mathrm{x}),\mathrm{t}),\\ Happens(\mathrm{e}'(\mathrm{x}),\mathrm{t}),\ t < now,\ \neg Happens(\mathrm{e}'(\mathrm{x}),now)\ \texttt{succeeds}$$

The query can be further reduced by adding the equalities $\mathrm{f} = \mathrm{s}$, $\mathrm{x} = \mathrm{y}$, and we finally obtain[24]

(28) $$?HoldsAt(\mathrm{h},\mathrm{t}),\ HoldsAt(\mathrm{f}(\mathrm{x}),\mathrm{t}),\ Happens(\mathrm{e}(\mathrm{x}),\mathrm{t}),\\ Happens(\mathrm{e}'(\mathrm{x}),\mathrm{t}),\ t < now,\ \neg Happens(\mathrm{e}(\mathrm{x}),now)\ \texttt{succeeds}$$

Now consider the DRS for the preferred reading of (21) the one in which the pronoun *She* refers back to *A delegate*. This reading results from unifying the variables x and y which is allowed since x is accessible for y.

(29) $$\boxed{\begin{array}{cccc} x & t & e & e' \\ delegate(x) & t < n & e \subseteq t & e' \subseteq t \\ & e : arrive(x) & e' : register(x) & \end{array}}$$

The integrity constraint for this DRS is precisely the one obtained in (28). The reader may obtain a clearer picture of what is going on here by rephrasing the preceding considerations as an inference problem. Recall from section 2.2 that an important goal of the formalism presented here

---

[23] For this reason we do not account for the preferred temporal ordering of the events, in which the registration of the delegate takes place *after* her arrival. The derivation of this effect proceeds along the lines sketched for the French examples in section 3.4. For a fuller treatment the reader is referred to van Lambalgen and Hamm [2004], especially Chapter 9.

[24] Unification applied to an integrity constraint is always hypothetical: it may be impossible to satisfy the query after the unification has been applied. Indeed we shall shortly see an example where this is so.

is to obtain an inference mechanism applying to DRSs. Now, taken in isolation, sentences (21) license the inference that a delegate registered[25]. This inference is non-monotonic, however, since it no longer holds if the premise set is enlarged to

(30) A delegate arrived. His wife arrived somewhat later. She registered (as accompanying person).

In this case it should no longer follow that a delegate registered. Formally, this failure is captured by the fact that the unification $f = s$, $x = y$ makes the query in (27) (or rather its extension to (30)) unsatisfiable, since y will be forced to be male, whereas x must be female. This observation can be recast in the form of an inference relation on integrity constraints, which automatically extends to DRSs:

Definition 1. *Let an argument with premises* Γ *and conclusion* φ *be given. Suppose* Γ *corresponds to the integrity constraint* $?G_0$ `succeeds`*, and* φ *corresponds to the integrity constraint* $?G_1$ `succeeds`*. Then* φ follows from Γ *if any update satisfying* $?G_0$ *can be extended to an update satisfying* $?G_1$*. Since DRSs can be made to correspond to integrity constraints, the same characterization applies to DRSs.*

4.1.3. Computational incorporation of lexical meaning

For DRSs containing complex DRS conditions conversion into integrity constraints is, we noted, not so straightforward. But at least in some instances these integrity constraints can be found. Our last example is an illustration of this general point.

The sentences in (31) are both ambiguous between a reading in which *on Monday* has scope over *often* (so that the phone calls in question all took place on some particular Monday) and one in which *often* has scope

---

[25] We do not claim that this approach provides a *general* theory of anaphora resolution. Such a claim would certainly be premature, since there are many different types of anaphora and moreover anaphora resolution is highly language dependent. Even typologically closely related languages such as English and German employ different strategies for anaphora resolution.

over *on Monday* (so that the phone calls took place at a number of different Mondays).[26]

(31) a. Max often called on Monday.
b. On Monday Max often called.

One of the problems that sentences like these present for semantics is exactly how they should be represented, and a second problem is how they are disambiguated – if and when they are – in context.[27] Here we concentrate on the simpler one of the two representations, in which *on Monday* has scope over *often*. In this case the Monday referred to by *on Monday* ist the period within which Max often called. (32) represents this reading in (somewhat simplified) DRT format.

(32)

$$\boxed{\begin{array}{c} n\ m\ t_m \\ t_m < n \\ Monday(t_m) \\ Max(m) \\ \boxed{\begin{array}{c} t^q \\ t^q \subseteq t_m \end{array}} \quad \left\langle \begin{array}{c} often \\ t^q \end{array} \right\rangle \quad \boxed{\begin{array}{c} e \\ e : call(m) \\ e \subseteq t_q \end{array}} \end{array}}$$

To convert this DRS into an integrity constraint we need to see how the meanings of, in particular, *often* and *on Monday* can be expressed within the formalism. We first treat *often*. Here is one way to capture its semantics. We assume that the meaning of *often* can be paraphrased as: the number of X's satisfying a given condition C exceeds a certain contextually determined limit (intuitively: the number of X's satisfying C that might have been expected). In order to formalise this intuition we use two built-in predicates of the logic programming language Prolog which we will now explain. The first is the three-place predicate *setof*: $setof(\mathrm{S}, \mathrm{C}, \mathrm{X})$ means that S is the set of X which satisfy condition C; in

---

[26] There may be a tendency for the second reading in the case of (31-a) and for the first in the case of (31-b), but to our judgement both sentences are definitely ambiguous.

[27] A more detailed discussion of (31) can be found in Reyle, Rossdeutscher and Kamp [in press].

set theoretic notation $S = \{X \mid C(X)\}$. The second is the two-place predicate $\mathit{length}(L, Y)$ which means that number Y is the length of list L. Let us first demonstrate how the adverbial quantifier *often* expressed by these Prolog–predicates works for the following simpler example involving past tense:

(33) Max often called.

The formalisation of (33) is given by integrity constraint (34). In order to faciliate reading we write the (untensed) event type "Max-call" (formally $\exists t.\mathit{call}[m, t]$) as e. As in our previous examples, the context is provided by the fluent f.

(34) $?\mathit{HoldsAt}(f, s),\quad s < \mathit{now},\quad \mathit{length}(S, y),\quad \mathit{setof}(S, \{\mathit{HoldsAt}(f, t), \mathit{Happens}(e, t), t < \mathit{now}\}, t),\ y \geq N$ `succeeds`.

The number N in (34) is the contextually expected number of times satisfying condition $\mathit{HoldsAt}(f, t)$, $\mathit{Happens}(e, t)$, $t < \mathit{now}$. *Often* then says via $\mathit{length}(S, y)$ and $y \geq N$ that the number of instances[28] satisfying condition $\mathit{HoldsAt}(f, t)$, $\mathit{Happens}(e, t)$, $t < \mathit{now}$ is greater than the contextually given expected number N.

We still have to analyze the phrase *on Monday*. For this we assume that we are given programs which specify seven fluents $f_{Su}, f_{Mo}, \ldots, f_{Sa}$ which together partition the time line in days of the week. We can then write a program clause which defines the notion of closest Monday f_{CMo}, i.e. last Monday or coming Monday.

(35) $\mathit{HoldAt}(f_{Mo}, s) \wedge |\mathit{now} - s| \leq 7\ \mathit{days} \rightarrow \mathit{HoldsAt}(f_{CMo}, s)$

With these definitions we are able to write integrity constraints for the DRS in (32). It is clear by now how to use the discourse referents *n* and *m*. The duplex condition for *often* introduces a variable y and the following integrity constraint.

$\mathit{length}(S, y),\ \mathit{setof}(S, \{\mathit{HoldsAt}(f, t), \mathit{Happens}(e, t)\}, t),\ y \geq N$ `succeeds`

The variables f and e are unified with material in the restrictor and nuclear scope: the fluent f_{Mo} corresponding to t^m in the restrictor and the event type a corresponding to *anrufen* in the nuclear scope. The result is:

---

[28] For simplicity we assume here that the tokens are points, but it is possible to generalize this to intervals.

$?length(\mathrm{S},\mathrm{y}),\ setof(\mathrm{S},\{HoldsAt(\mathrm{f}_{\mathrm{Mo}},\mathrm{t}),Happens(\mathrm{a},\mathrm{t})\},\mathrm{t}),$
$\mathrm{y} \geq \mathrm{N}$ `succeeds`

The condition $t^m < n$ adds further information about the temporal location of event and reference time. We thus get:

$$?HoldsAt(\mathrm{f}',\mathrm{s}),\ \mathrm{s} < now,\ length(\mathrm{S},\mathrm{y}),\ setof(\mathrm{S},\{HoldsAt(\mathrm{f}_{\mathrm{Mo}},\mathrm{t}),Happens(\mathrm{a},\mathrm{t}),\mathrm{t} < now\},\mathrm{t}),\ \mathrm{y} \geq \mathrm{N}\ \texttt{succeeds}$$

The additional condition $Monday(t^m)$ in the upper DRS has the effect of unifying the terms f' and f_{Mo} with $\mathrm{f}_{\mathrm{CMo}}$. The resulting integrity constraint for DRS (32) therefore is (36).

(36) $?HoldsAt(\mathrm{f}_{\mathrm{CMo}},\mathrm{s}),\ \mathrm{s} < now,\ length(\mathrm{S},\mathrm{y}),$
$setof(\mathrm{S},\{HoldsAt(\mathrm{f}_{\mathrm{CMo}},\mathrm{t}),Happens(\mathrm{a},\mathrm{t}),\mathrm{t} < now\},\mathrm{t}),\ \mathrm{y} \geq \mathrm{N}$
`succeeds`.

Note that condition s < *now* automatically picks out last Monday instead of coming Monday. The events quantified over by *often* thus take place entirely in the past. The translation of sentence (31-a) with wide scope reading for *often* proceeds analogously and is left to the interested reader.

5. Concluding remarks

Let us reiterate what we see as the main points of this article. We emphasized the necessity of a computational approach to semantics, if it wants to establish a truly productive interaction with cognitive (neuro)science. We singled out computations of event structures and discourse models for particular attention. It was argued that these computations can be viewed as identical in structure to those executed by the human planning mechanism, thus leading to the conjecture that in the course of human evolution the planning system was co-opted for purposes of language comprehension. We substantiated these claims by sketching a formalism consisting of an event calculus, a planning formalism from robotics, together with a nonmonotonic inference engine, constraint logic programming. We showed informally how computations of discourse models could proceed, and we discussed an ongoing experimental investigation using ERP which attempts to find traces of these computations.

The computational approach to natural language semantics, to which this article has been primarily devoted, shares most of its basic goals with DRT. Because of its computational advantages the approach could be seen as a wholsesale alternative to DRT, which could replace it altogether. But at the present time it is too early to assess whether this is a genuine possibility. First, we do not yet know whether all of the information that natural languages can express (and which DRT has made it a primary concern to develop the formal means to represent) can be adequately expressed within the formalism proposed. Second, even if this were possible, DRSs might still be useful as intermediaries between natural language and representations in the event calculus, allowing the DRS construction algorithm to deal with all kinds of features of the syntax-semantics interface which the proposed formalism is not naturally equipped to deal with. And for all we can tell at present this intermediate level may not be just a convenience for semantic theory, but a cognitive reality, no less than the event calculus-based representations which have been the central focus of this paper.

References

Comrie, B. (1976). *Aspect*. Cambridge University Press, Cambridge.

Elworthy, D. (1992). *The Semantics of Noun Phrase Anaphora*. PhD thesis, Cambridge University.

—(1995). A theory of anaphoric information. *Linguistics and Philosophy*, 18.

Foster, C. (1992). *Algorithms, Abstraction and Implementation*. Academic Press, London.

Goldberg, A. E. (1995). *Constructions: a construction grammar approach to argument structure*. University of Chicago Press, Chicago, Ill.

Groenendijk, J. and Stokhof, M. (1990). Dynamic montague grammar. In Kálmán, L. and Pólós, L. editors, *Papers from the second Symposium on Logic and Language*. Budapest.

—(1991). Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–100.

Hagoort, P., Hald, L., Bastiaansen, M. and Petersson, K. M. (2004). Integration of word meaning and world knowledge in language comprehension. *Science*, 304:438–441.

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts at Amherst.

Jackendoff, R. (1994). *Consciousness and the Computational Mind*. MIT Press, Boston.

—(1998). Why a conceptualist view of reference? a reply to abbott. *Linguistics and Philosophy*, 21:211–219.

—(2002). *Foundations of Language*. Oxford University Press, Oxford.

Kamp, H. (1981). A theory of truth and semantic representation. In Groenendijk, J., Janssen, T. and Stokhof, M. editors, *Formal Methods in the Study of Language*. Amsterdam: Mathematical Centre Tracts 135, Amsterdam.

Kamp, H. and Reyle, U. (1993). *From Discourse to Logic*. Reidel, Dordrecht.

—(1996). A calculus for first order discourse representation structures. *Journal of Logic, Language and Information*, 5.

Koechlin, E., Basso, G., Pietrini, P., Panzer, S. and Grafman, J. (1999). The role of the anterior prefrontal cortex in human cognition. *Nature*, 399:148–151.

Krifka, M. (1996). Parametrized sum individuals for plural reference and partitive quantification. *Linguistics and Philosophy*, 19.

Lewis, D. (1972). General semantics. In Davidson, D. and Harman, G. editors, *Semantics of Natural Language*. Reidel, Dordrecht.

Nouwen, R. (2003). *Plural Pronominal Anaphora in Context: Dynamic Aspects of Quantification*. PhD thesis, Utrecht University.

Reyle, U., Rossdeutscher, A. and Kamp, H. (2005). Ups and downs in the theory of temporal reference. University of Stuttgart.

Rossdeutscher, A. and Reyle, U. (2000). Constraint-based bottom up discourse interpretation. In Reyle, U. editor, *Presuppositions and Underspecification in the Computation of Temporal and other Relations in Discourse*. Arbeitsberichte des Sonderforschungsbereiches 340, N. 157, Stuttgart/Tübingen.

Steedman, M. (1997). Temporality. In van Benthem, J. and ter Meulen, A. editors, *Handbook of Logic and Language*. Elsevier, Amsterdam.

Stenning, K. and van Lambalgen, M. (2005). Semantic interpretation as reasoning in nonmonotonic logic. to appear in Cognitive Science.

Suddendorf, T. and Corballis, M. (1997). Mental time travel and the evolution of mind. *Genetic, social and general psychology monographs*, 123:133–167.

Tomasello, M. (2003). *Constructing a language. A usage-based theory of language acquisition*. Harvard University Press, Boston.

Trabasso, T. and Stein, N. L. (1994). Using goal-plan knowledge to merge the past with the present and future in narrating events on line. In Haith, M. H., Benson, J. B., Roberts, R. J. and Pennington, B. F. editors, *The development of future-oriented processes*, pages 323–352. University of Chicago Press.

van den Berg, M. (1996). Some Aspects of the Internal Structure of Discourse: The Dynamics of Nominal Anaphora. PhD thesis, Amsterdam University.

van Lambalgen, M. and Hamm, F. (2004). *The proper treatment of events*. Blackwell, Oxford and Boston.

Verhagen, A. (2005). *Constructions of intersubjectivity*. Oxford University Press, Oxford.

von Genabith, J., Kamp, H. and Reyle, U. (2005). Discourse representation theory. University of Stuttgart, to appear in: Gabbay, D. (ed.), Handbook of Philosophical Logic.