TEMPORAL REASONING IN NATURAL LANGUAGE UNDERSTANDING: THE TEMPORAL STRUCTURE OF THE NARRATIVE

Alexander Nakhimovsky
Department of Computer Science
Colgate University
Hamilton, NY 13346 USA
CSNet: sasha@colgate

Abstract

This paper proposes a new framework for discourse analysis, in the spirit of Grosz and Sidner (1986), Webber (1987a,b) but differentiated with respect to the type or *genre* of discourse. It is argued that different genres call for different representations and processing strategies; particularly important is the distinction between subjective, performative discourse and objective discourse, of which narrative is a primary example. This paper concentrates on narratives and introduces the notions of temporal focus (proposed also in Webber (1987b)) and narrative move. The processing tasks involved in reconstructing the temporal structure of a narrative (Webber's e/s structure) are formulated in terms of these two notions. The remainder of the paper analyzes the durational and aspectual knowledge needed for those tasks. Distinctions are established between grammatical aspect, aspectual class and the aspectual perspective of a sentence in discourse; it is shown that in English, grammatical aspect under-determines the aspectual perspective.

NARRATIVES

This paper investigates the varieties of temporal knowledge and temporal reasoning that are at work in understanding extended narratives. It starts out by developing a new framework for narrative representation, a framework that has developed independently from, but is very similar to Webber, 1987a, 1987b. It also builds on the ideas of Grosz and Sidner (1986), but reformulates them specifically for the task of narrative understanding. A reformulation, I believe, is needed because different genres of discourse - narrative, expository text, task-oriented dialog, argument, etc. - have different principles of organization that call for different representations

and processing strategies. Without offering a comprehensive taxonomy of discourse genres I would like to stress that narrative stands out by virtue of its two properties: it is objective and it unfolds in time.

A distinction between subjective and objective modes of discourse has been drawn by many authors in linguistics and structuralist poetics, who all "have a category of narration to which another category is opposed; and they all agree that the non-narrative category is more subjective" (Lyons, 1982:117). One manifestation of the objectivity of narratives is the structure of the underlying intentions. This structure plays an important role in Grosz and Sidner, 1986 who propose, *inter alia*, are that (a) the content of discourse is embedded in, and classified by, the speaker's intentions which form a hierarchical intentional structure, and (b) the content structure is separate from the attentional state, and both are rather indirectly represented by the linguistic material of discourse, organized in a hierarchical structure of discourse segments. I adopt (b) without reservations, but (a), I suggest, needs to be modified and differentiated. In dialogs the structure of intentions is, indeed, rich and informative (note that most indirect speech acts occur in dialogs); in narratives and expository prose the intention is practically constant: "intend that the other discourse participant believe proposition p" (cf. Grosz and Sidner, 1986:184). In other words, the only discourse purpose of a narrative or its segments is to modify the memory of the other discourse participant. Removing this, rather uninformative, top level of intention, reveals the "objective" content structure of the narrative, whose main building block is a situation persisting or evolving in time, best visualized as a four-dimensional piece of time-space. Loosely following Hayes, 1978 I use the term history-token (h-token) for all varieties of such situations (events, processes, activities, ha-

bitual actions, etc); each h-token is an instance of a history-type (h-type) corresponding to abstract situations types of Situation Semantics. I assume that associated with each predicate of the meaning representation language is a set of roles such as Agent, Object or Patient; an h-type is a predicate together with its roles and a selectional restriction on them (cf. Creary and Pollard, 1985, Hobbs et al, 1986).

Removing the top layer of intentions leads to other changes in the Grosz-Sidner model. Each discourse segment (DS) is now characterized by its main h-token, rather than its DS purpose. An h-token is, in turn, characterized by a spatio-temporal location, a set of participants and a time scale. Dominance relations between intentions correspond to compositional relations between h-tokens: the h-token of entering a room decomposes into opening the door, crossing the threshold, closing the door (provided there is a door to open and close). Satisfaction-precedence relations between intentions correspond to the temporal and causal relations between histories. Thus re-interpreted, the pair intentional structure-attentional state of Grosz and Sidner, 1986 becomes very similar to Webber's (1987a:137) proposal: "Along with building up a discourse model of the entities salient to the given text, the listener is also building up a model of the events and situations they participate in-e/s structure." (Although Webber speaks of a 'text' in general, I believe she means 'a narrative text,' and all her examples are such.) To emphasize the similarity of the two approaches, and to avoid proliferation of terminology, I use Webber's term *e/s structure* for the representation of the narrative's content, but retain Grosz and Sidner's terminology for the attentional state and speak of a focus space (FS) corresponding to each DS, and a focus space stack (FS stack). An important difference is that I don't think anything ever gets popped off the FS stack: it just keeps growing, representing the linear progression of the text (while the *e/s structure* represents the temporal progression of its content). It is a stack only in the sense that its top element is the easiest to access, not in the sense of following the LIFO discipline. Even interruptions, digressions and flashbacks, to which the pop-off action seems most applicable, are better represented as a move into a new FS, accompanied by a promise to return: to return to the immediately preceding FS in the case of interruptions, and to a specified position in the *e/s structure* in the case of digressions and flashbacks.

The constancy of intention is one aspect of the narrative's objectivity; another one is its "closeness unto itself" in the processing of definite and temporal anaphora. Subjectivity goes with deixis, the constant presence of the situation of utterance in the processing model. Objective texts' contents are removed from deixis into a separate universe, which, in the case of narratives, is endowed with its own, separate timeline. In some languages this separateness is clearly signalled by special narrative-beginning devices and/or narrative tenses (Dahl, 1985). In English, there is of course an overlap between the "narrative" and "non-narrative" tenses, but it is far less complete than is usually supposed: one could go through a book on computer science and not find a single occurrence of a past tense, except, perhaps, in short passages on the history of particular ideas; conversely, one could go through a long novel and not find a single sentence in the present or future, except in the characters' dialogs.

Behind the superficial difference in the use of tenses stands the more important one in the basic meaning of the grammatical category of tense. The standard view is that tense indicates relative position in time with respect to the speech event (Comrie, 1985). In dialogs tense indeed appears in its deictic function, which is also the dominant function of the present and future tenses. However, past tenses are different, especially in narratives; consider: "On March 5, 3275, Captain Kirk got up early, shaved and boarded the Enterprise." Surely, the form of the verb *shave* does not mean that the Captain was clean-shaven before the book went to print. Rather, it indicates that we are in a narrative, and it helps position the event vis-a-vis the narrative's preceding events. In other words, narrative tenses are anaphoric, not deictic. An analogy with pronouns is, perhaps, useful: although 3 person pronouns are grouped together with *I* and *you* in traditional grammars, and although they can be used deictically (if strongly accented and accompanied by a gesture) their primary function is anaphoric.

The anaphoric nature of past tenses (first recognized in Partee (1973), investigated specifically in narratives in Hinrichs (1986)) has important computational implications, for anaphora can only be resolved with respect to a constantly maintained and updated focus (Grosz, 1977; Sidner, 1983). To emphasize the parallel between temporal and definite anaphora, I will speak of the temporal focus of a narrative. (The same term for the same concept and

with the same motivation is proposed in Webber, 1987b; in Nakhimovsky 1986, 1987 I speak of the Active Window on discourse, or Window for short; Kamp and Rohrer, 1983 have recycled Reichenbach's Reference Point for a similar concept.) If the focus *simpliciter* answers the question "What are we talking about?" the temporal focus answers the question "Where in time is the narrative now?" As the narrative progresses, the temporal focus changes its position in time; I will refer to the movement of temporal focus from one sentence of the narrative to the next as *the narrative move*.

A narrative move can remain within the current FS, or shift to a different one, which can be totally new or a resumption of an old FS from the stack. (In terms of linguistic structure, the current sentence may continue the same, or start a new, DS.) The two kinds of narrative moves will be called micro- and macro-moves, respectively. Examples (1)-(3) contrast the two kinds of moves and illustrate other concepts introduced in this section.

(1) a. John entered the president's office. b. The president got up.

This is narrative at its simplest: an orderly progression of events within the same narrative unit. The required inferential work is relatively transparent. The event of John's entering the office results in the state of his being in the office: this is part of the lexical meaning of enter. The temporal focus is inside this state, at its beginning. Sentence b., which in isolation could mean that the president got up from his bed at home, is interpreted *vis-a-vis* the position of the temporal focus: the president was in his office, sitting; he saw John and got up; both men are now standing, 'now' referring to the temporal focus as it always does. This example shows that it would be more accurate to speak of the spatio-temporal focus to which the current situation is anchored (cf. Barwise and Perry, 1983) but I leave the spatial dimensions of narrative for future research.

Examples (2) and (3) illustrate macro-moves:

(2) a. Gradually, Harvey began to yield the details of his crime, prodded by the persistent questions of the investigator. b. He arrived at the bank at 4 p.m. dressed as a postal worker.

(3) a. Hartley and Phoebe had been sent by their mother to fix the tail valve of the windmill. b. In the great expanse of the prairie where they lived, the high tower of the windmill was the only real landmark (Worline, 1956:1).

In (2), the similarity between definite and temporal anaphora stands out quite clearly. Just as *he* in sentence b. anaphorically evokes discourse-prominent Harvey, so *arrived* evokes the time of the discourse-prominent crime event and 4 p.m. evokes the day of that event. Just as *he* selects for anaphoric reference one of two discourse entities available for pronominalization, so *arrived* and 4 p.m. select one of two available events, the interrogation and the crime. The shift of temporal focus to an earlier event, over a considerable time interval, signals the beginning of a new DS. The FS associated with the old DS is saved on the stack together with the last position of the temporal focus in it, which is under-determined by the English narrative: it can be within, or right after, the reconstructed details history. If the DS is resumed with *Harvey took a sip of water and mopped his brow*, we don't know whether the reconstruction is over or not.

In (3) the beginning of a new DS in sentence b. is indicated by a drastic change in time scale, rather than movement of focus. Sentence a. establishes, either directly or through simple, lexicon-based inferences, three events: the tail vane broke, mother sent the children to fix it, the children set off walking. The temporal focus, indicated by the past perfect tense, is in the middle of the walking event; the time scale of the entire sequence is within a day or two. The time scale of sentence b, indicated by the *where they lived* clause and the lifetime of a windmill (McDermott, 1982), is years or decades. (Note the accompanying shift in the spatial scale from one household to the entire prairie.)

Narratives (1)-(3) illustrate several important points about the temporal focus. First, it is always inside some history, either directly narrated or inferred. If that history has a built-in terminal point that is reached in the normal course of events, the position of the focus sets up the expectation that, within a certain time scale, the terminal point will be reached. So, in (3) we expect the children to make it to the windmill before it gets dark, and indeed, after a page of background material, the FS of (3a) is resumed, with children already standing at their destination. Second, the position of the temporal focus may be under-determined, as in (2), but there are precisely two possibilities: inside or right after the most recently narrated history. Adopting the terminology of Smith (1986) I will speak of the imperfective and perfective sentence perspective, respectively.

Given the conceptual apparatus that has

been developed in this section, several tasks involved in narrative understanding can be specified. The tasks are clearly interrelated, but in this paper I make no comment on how the interaction can be set up.

(4) As each new sentence of the narrative comes in do:

- a. *determine the type of narrative move (micro or macro) that the new sentence represents. If it is a macro-move, update the FS stack and position the new FS in the existing e-s structure. If it is a micro-move, determine the temporal relations between the histories described by the current and the preceding sentence.*
- b. *using knowledge about durations and aspectual classes of events, determine the aspectual perspective of the new sentence and the position of the temporal focus;*
- c. *using knowledge about causality and internal constituency of events, add inferred events to the narrated ones; update old expectations and set up new ones.*

Several kinds of temporal knowledge are thus brought to bear on the process of narrative understanding. First, there is knowledge about durations and time scales, and the interaction, totally disregarded in existing work, between the event structure of the narrative and the hierarchy of "received" time cycles such as times of day, seasons of the year and the stages of human life. Second, there is compositional knowledge about internal constituency of events and their terminal points. Third, there is aspectual knowledge, both lexical, about intrinsic properties of histories, and grammatical, about the way the history is presented by a given verb form. The remainder of this paper investigates these three kinds of knowledge and the ways they are represented in the lexicon and utilized in narrative understanding.

DURATION

Information about durations can be entered in the lexicon in the following three ways that are not mutually exclusive: (a) most generally, as qualitative functional dependencies (Forbus, 1985) among the participants of the situation; so, the time it takes to read a text depends on its length and genre, and the proficiency of the reader; (b) for some h-types (e.g. lecture,

shower, lunch) the duration of their h-tokens is stable and can be entered in the lexicon directly as a fuzzy number (e.g. lecture [1,2 hour]; (c) for a majority of h-types, the time scale of their h-tokens is quite narrowly constrained, where the time scale of an interval is a sequence of measurement units that are "natural" to it: measured in a natural unit, the length of the interval will not be a very small fraction (greater than some constant R) or a very big number (less than some constant N). The important ideas are, first, that measurement units form a small set that is partially civilization specific, partially determined by the biological and physical universals; second, that the duration of an h-token constrains the choice of measurement units in which its duration is measured and thus the precision of measurements: when we say *It took John an hour to repair a faucet* we don't mean that it took him 3600 seconds.

An important durational class of h-tokens is instantaneous events. There is a persistent misconception, inspired by scientific thinking, that the notion of an instantaneous or punctual event can only be defined relative to a time scale because "we can always 'increase the magnification' and find more structure" (Allen and Kautz, 1985:253; see also Dowty, 1986, Kamp, 1979). I believe that instantaneousness is an absolute quality determined by our biology: instantaneous events are those that are not perceived by humans as possessing internal structure. Languages select such events for special treatment by disallowing the "imperfective description" of them: one cannot use the imperfective aspect to place the temporal focus in the middle of an instantaneous event, so that *The light was flashing* does not place the temporal focus inside an individual flash. (More on aspects below.)

Non-instantaneous events are, intuitively, discrete and countable entities with a distinct beginning and end; packaged in between the beginning and end of an event is the "stuff the event is made of," which is a process or state. This intuition is discussed in a considerable body of literature that compares the event-process and count-mass oppositions (Mourelatos, 1981, Bunt, 1985, Bach, 1986). As I argue in Nakhimovsky (1986), all these authors should also have allowed for events made out of states, as, for example, the event described by *Bobby took a nap*. Surprisingly, collocations of this nature have never, to my knowledge, been discussed in connection with the English aspectual system. (Cf. also *did some reading, went*

for a walk.)

The distinctions event-process and process-state are thus orthogonal to each other, rather than forming a single classification as in Mourelatos, 1981; Allen, 1984. The former distinction is one of aspect: "the term 'process' means a dynamic situation viewed imperfectly, and the term 'event' means a dynamic situation viewed perfectly" (Comrie, 1976:51). The latter distinction is one of aspectual class. This is elaborated in the next section.

ASPECT

In what follows it is essential to keep the following three concepts apart: aspect as a grammatical category of the verb, implemented by affixes, auxiliaries and such; aspectual class, which is a characteristic of an h-type or lexical meaning; the aspectual perspective of the sentence. Both grammatical aspect and aspectual class sometimes uniquely determine, sometimes just strongly constrain, the aspectual perspective. In English, the progressive aspect guarantees that the sentence perspective is imperfective; in any language, instantaneous events are presented perfectly (which does not mean that the corresponding verbs are in any sense perfective). All three concepts are needed for understanding the workings of aspectual systems; I don't think anybody in the abundant recent literature on aspect keeps all three clearly apart.

There are languages, most notably Slavic, where the difference in the sentence perspective is hard-wired into verb morphology: simplifying slightly, every Russian verb is either perfective or imperfective, and the morphological feature of the verb determines the aspectual perspective of the sentence. (In fact, the English term 'aspect' is a mistranslation of the Russian term 'vid,' 'view, perspective.') In other words, I claim, rather audaciously, that grammatical aspect is a purely attentional device that helps determine the position of the temporal focus; all the other shades of aspectual meaning result from interactions between this (pragmatically defined) *Grundbedeutung* and numerous other factors, including aspectual class, discourse genre, and general pragmatic principles of language.

The following examples, adopted from Dowty (1986), illustrate the interplay between aspect, aspectual class and the micro-move of the narrative. (I repeat (1) here for convenience.)

(1) a. John entered the president's office. b. The president got up.

(5) a. John entered the president's office. b. The president was asleep. c. The clock on the wall ticked loudly.

(6) a. John entered the president's office. b. The president was writing a letter.

Sentences (1a) and (1b) describe two processes (entering and getting up) that each have a built-in terminal point that is reached in the normal course of events and beyond which the processes cannot continue. (In Vendler's (1967) well-known classification such processes are called accomplishments; I call them, following Comrie (1976), telic processes.) The aspectual perspective of both sentences is perfective; the events of the two sentences are understood to have happened in succession; the temporal focus has advanced to the time when both men are standing.

Sentences b. and c. in (5) describe a state and an atelic process, respectively. They are understood to have begun before the event of sentence 1, and to persist in parallel. The temporal focus stands still. Note that the sentence perspective of b. and c. is determined by the aspectual class, not grammatical aspect. In (6), however, the sentence perspective of b., and the micro-move from a. to b., are determined by the progressive form of the verb: although writing a letter is a telic process the micro-move in (6) is the same as in (5).

The history of misconceptions concerning the English aspectual system can be summarized as follows. First it was believed that English has no aspect; progressive was called a tense. When it came to be recognized that progressive is a variety of the imperfective aspect, the next misconception was to assume that since English has an imperfective, it ought to have a perfective also, with simple past an obvious candidate. However, examples like (5c) show that a sentence with a verb in simple past can have the imperfective perspective. The current consensus seems to be that simple past of accomplishment verbs is perfective (Hinrichs, 1986:68; Dowty, 1986:46-8). In other words, if the verb form = simple past and the aspectual class = telic process then the sentence perspective is perfective and the temporal focus advances. Consider, however, example (7), where two accomplishments, both described by verbs in the simple past, unfold in parallel and are both interrupted by a doorbell:

(7) a. After supper, Alice and Sharon sat down in the living room. b. Alice read a book, Sharon watched her favorite ballet on television. c. Suddenly the doorbell rang.

Other examples of micro-moves that violate Hinrichs' rule are given in (8) and (9), quoted from Dowty, 1986. (The rule can also be violated by a macro-move, as in example (2)).

(8) John knelt at the edge of the stream and washed his hands and face. He washed slowly, feeling the welcome sensation of the icy water on his parched skin. (From Dry, 1983)

(9) Pedro dined at Madam Gilbert's. First he gorged himself on hors d'oeuvres. Then he paid tribute to the fish. After that the butler brought in a glazed chicken. The repast ended with a flaming dessert. (From Kamp, ms.)

I conclude that English has no (morphological) perfective; it has a marked imperfective and an unmarked default that does not provide substantial information about the aspectual perspective of the sentence (cf. Dahl, 1985 for the same view). In other words, English morphology, even combined with aspectual class, underdetermines the sentence perspective and the micro-move of the narrative. However, the number of possibilities is limited, and an extensive empirical investigation could, I believe, produce a full catalog of micro-moves commonly employed in Western narratives.

ASPECTUAL CLASS

The major division among non-instantaneous histories, recognized at least since Aristotle, is between process (*energeia*) and state (*stasis*). In recent times, Vendler (1967) proposed a highly influential classification that is still commonly accepted, although the principles of classification have changed. Vendler believed, erroneously, that he was classifying English verbs, rather than sentence denotations, and he used such language-specific criteria as whether or not a verb has a progressive form (Vendler's statives, such as *know*, don't). In the model-theoretical version of Taylor and Dowty, the classification is based on the relationship between the truth value of a sentence at an interval and at its subintervals; so, for instance, a sentence *S* is stative (denotes a state) iff it follows from the truth of *S* at an interval *I* that *S* is true at all subintervals of *I*. (Dowty, 1986:42).

I submit that these criteria cannot possibly be right, i.e. capture the real distinctions operative in the workings of human language: these have to relate to something perceived and experienced, rather than truth values (which is not to deny that real distinctions may result in fairly consistent truth-functional properties). It is not accidental that Dowty's own example of a state (sleep) contradicts his definition: we can truthfully say that Bob slept from 10 to 6 even if he got up once to go to the bathroom. My proposal is that we take the physical vocabulary of processes and states seriously, and classify histories according to their internal dynamics, the stability of their parameters and the resources they consume. (Part of the internal dynamics, in the presence of a conscious agent, is the degree of volitional control.) We can then note the distinction between states that do not require any resources to sustain themselves (know English, own a house) and those that do (sleep requires a certain amount of sleepiness that gradually wears out). The sub-interval property holds only for zero-resource, zero-control states, and is, in fact, a simple consequence of their other properties: a state that requires no resources and cannot be dropped in and out of at will obtains continuously.

Resource-consuming states all seem to require only generic internal resources, which are not specific to any given state but rather to all individuals of a given sort. Within processes, there are those that require only generic resources (walking) and those that require process-specific resources as well: reading, for example, requires not only being awake and not too hungry, but also a text to read. Telic processes can be defined as processes that consume a specific amount of a domain-specific resource. Resources are understood broadly: walking from A to B consumes the distance between them, building a house consumes the as-yet-unbuilt but envisioned part of it, and destroying a house consumes the finite amount of "structure" or "order" built into it. These examples illustrate three main classes of telic processes: creating an object, destroying an object, and moving a specified amount of material (possibly the mover himself) to a specified destination. A subclass of destruction processes are ingestions, which convert an external resource into an internal one. Moving is understood to include all three of Schank's PTRANS, ATRANS and MTRANS classes, with the proviso that, unlike physical motion, MTRANS really copies structures from the source to the destination. Moving also includes gradual (but not

instantaneous) changes of state.

Lacking internal structure, instantaneous events have to be classified by comparing the world before and after them. An instantaneous event can terminate either a process or a state, and it can initiate either a process or a state; if it is sandwiched in between two processes or two states, the two can be the same or different. The resulting classification, discussed in Nakhimovsky, 1987, captures linguistically significant distinctions: for instance, most English verbs describing instantaneous events fall into those groups where the instantaneous event meets a state.

FUTURE RESEARCH

Perhaps the biggest task involved in narrative understanding is to infer, using knowledge about causality and the internal constituency of events, the missing links between narrated events and the temporal relations between them. This involves solving qualitative functional equations that hold between the parameters of described histories and resources they consume (cf. Forbus, 1985), and propagating durational constraints (cf. Allen and Kautz, 1985). An analysis of the required lexical knowledge is presented in this paper and Nakhimovsky (1987). The subject is further developed in Nakhimovsky (in preparation).

ACKNOWLEDGEMENTS

I'm grateful to Colgate University for giving me a leave to do research, and to Yale's AI Project for providing a stimulating environment for it. My conversations with Tom Myers, Donka Farkas and Larry Horn have helped me clarify my ideas. Alex Kass read a draft of the paper and provided valuable comments and invaluable technical assistance.

REFERENCES

- Allen, James, and Kautz, H. A. 1985. A model of naive temporal reasoning. In Hobbs and Moore, 1985.
- Bach, Emmon. 1986. The algebra of events. *Linguistics and Philosophy*, 9:5-16.
- Bobrow, Daniel. (Ed.) 1985. *Qualitative Reasoning about Physical Systems*. Cambridge, MA: MIT Press.
- Bunt, H. 1985. The formal representation of (quasi) continuous concepts. In Hobbs and Moore, 1985.
- Comrie, B. 1976. *Aspect*. Cambridge: Cambridge University Press.
- Comrie, B. 1985. *Tense*. Cambridge: Cambridge University Press.
- Creary, L., and Pollard, Carl. 1985. A computational semantics for natural language. *ACL-85*: 172-9.
- Dahl, O. 1985. *Tense and Aspect Systems*. Oxford: Basil Blackwell.
- Dowty, David. 1986. The effects of aspectual class on the temporal structure of discourse. In Dowty, 1986:37-61.
- Dowty, David. (ed.) 1986a. Tense and Aspect in Discourse. *Linguistics and Philosophy*, 9,1. (Special issue)
- Dry, Helen. 1983. The movement of narrative time. *Journal of Literary Semantics*, 12:19-53.
- Forbus, K. D. 1985. Qualitative process theory. In Bobrow, 1985.
- Grosz, Barbara, 1977. The representation and use of focus in a system for understanding dialogs. *IJCAI-77 Proceedings*, Los Altos, CA: Morgan Kaufmann.
- Grosz, Barbara, and Candace Sidner. 1986. Attention, intentions and the structure of discourse. *Computational Linguistics*, 12,3:175-204.
- Hayes, P. 1978. The naive physics manifesto. In D. Michie (Ed.), *Expert Systems in the Microelectronic Age*. Edinburgh, Scotland: Edinburgh University Press.
- Hinrichs, Erhard. 1986. Temporal anaphora in discourses of English. In Dowty, 1986a: 63-82.
- Hobbs, J. R., and R. C. Moore (Eds.). 1985. *Formal Theories of the Commonsense World*. Norwood, NJ: Ablex Publishing Corporation.
- Hobbs, J. R. et al. 1986. Commonsense metaphysics and lexical semantics. *ACL-86*: 231-40.
- Hopper, Paul. 1978. Aspect and foregrounding in discourse. In Talmy Givon, ed., *Discourse*

- and *Syntax*. (Syntax and Semantics, vol. 12.) New York: Academic Press.
- Hopper, Paul. 1982. (ed.) *Tense-aspect: between Semantics and Pragmatics*. Amsterdam: John Benjamins.
- Kamp, Hans. 1979. Events, instants and temporal reference. In R. Bauerle, U. Egli and Arnim von Stechow (eds.) *Semantics from Different Points of View*. Berlin: Springer.
- Kamp, Hans. ms. Discourse representation and temporal reference.
- Kamp, Hans, and Christian Roerer. 1983. Tense in texts. In R. Bauerle, C. Schwarze, Arnim von Stechow (eds.), *Meaning, Use and Interpretation in Language*. Berlin: De Gruyter.
- Lyons, John. 1982. Deixis and subjectivity. In R. Jarvella and W. Klein (eds.) *Speech, Place and Action*. Chichester: John Wiley, 101-124.
- McDermott, Drew. 1982. A temporal logic for reasoning about processes and plans. *Cognitive Science*, 1982, 6:101-155.
- Mourelatos, A. P. D. 1981. Events, processes, and states. In Tedeschi and Zaenen, 1981.
- Nakhimovsky, A. 1986. Temporal components of word meanings and their coding in the lexicon. Proceedings of the Conference on Advances in Lexicology, University of Waterloo.
- Nakhimovsky, A. 1987. The lexicon, grammatical categories and temporal reasoning. In *Artificial Intelligence and Simulation of Behavior*, Chichester: John Wiley.
- Nakhimovsky, A. In preparation. Tense, aspect and the temporal structure of the narrative.
- Partee, Barbara. 1973. Some structural analogies between tenses and pronouns in English. *Journal of Philosophy*, 70:601-9.
- Sidner, Candace. 1983. Focusing in the comprehension of definite anaphora. In M. Brady and R. Berwick eds., *Computational Models of Discourse*, Cambridge, MA: MIT Press.
- Smith, Carlota. 1986. A speaker-based approach to aspect. In Dowty, 1986a: 97-115.
- Vendler, Z. 1967. *Linguistics and Philosophy*. Ithaca, NY: Cornell University Press.
- Webber, Bonnie. 1987a. Event reference. In *TINLAP-3 Position Papers*. Las Cruces: New Mexico State University.
- Webber, Bonnie. 1987b. Two steps closer to event reference. TR MS-CIS-86-74, Department of Computer and Information Science, University of Pennsylvania.
- Worline, Bonnie Bess. 1956. *The Children Who Stayed Alone*. New York: Scholastic Book Services.