Annu. Rev. Psychol. 1997. 48:163–89

DISCOURSE COMPREHENSION

Arthur C. Graesser

Department of Psychology, The University of Memphis, Memphis, Tennessee 38152

Keith K. Millis

Department of Psychology, Northern Illinois University, DeKalb, Illinois 60115

Rolf A. Zwaan

Department of Psychology, Florida State University, Tallahassee, Florida 32306

KEY WORDS: discourse processing, inferences, reading, text comprehension

ABSTRACT

The field of discourse processing has dissected many of the levels of representation that are constructed when individuals read or listen to connected discourse. These levels include the surface code, the propositional textbase, the referential situation model, the communication context, and the discourse genre. Discourse psychologists have developed models that specify how these levels are mentally represented and how they are dynamically built during comprehension. This chapter focuses on the meaning representations that are constructed when adults read written text, such as literary stories, technical expository text, and experimenter-generated "textoids." Recent psychological models have attempted to account for the identification of referents of referring expressions (e.g. which person in the text does *she* refer to), the connection of explicit text segments, the establishment of local and global coherence, and the encoding of knowledge-based inferences.

CONTENTS

0084-6570/97/0201-0163$08.00

INTRODUCTION

A distinguished experimental psychologist approached one of us at a conference and grumbled "Why do you waste your time studying discourse? Why don't you study something fundamental, like perception, learning, memory, or eyelid conditioning in rabbits?" Shocked and bewildered, a witty insightful response was never delivered to the elderly gentleman. However, we have periodically imagined the perfect response over the years. A flip answer would be "Because it's there." This answer would be entirely correct, of course. Adults in this culture do spend most of their conscious life speaking, hearing, writing, and reading various forms of connected discourse. A more confrontational response would be "Because discourse *is* fundamental." This response would reflect what we believe. Discourse is what makes us human, what allows us to communicate ideas, facts, and feelings across time and space. Introductory texts in cognitive psychology should have a chapter on discourse, just as there are chapters on perception, memory, learning, problem solving, and language. The practical response would be "Because it's useful." The currency of psychology rises to the extent that discourse psychologists can improve reading, text design, complex learning, and social interaction.

There are more insightful reasons for investigating discourse processing as a primary object of inquiry. First, discourse processing has some unique properties that cannot be reduced to other subareas of psychology, such as psycho linguistics and memory. Connected discourse is more than language per se, and much more than a sequence of individual sentences. Comprehension cannot be reduced to problems of accessing and constructing memory representations. Second, discourse spans enough context to constrain the interpretation of input in a systematic fashion. As one colleague put it, a sentence out of context is always ambiguous, whereas a sentence in a discourse context is rarely ambiguous. Third, some genres of discourse, such as stories, are microcosms of events and experiences in the real world. Both stories and everyday experiences include people performing actions in pursuit of goals, events that present obstacles to these goals, conflicts between people, and emotional reactions. Experimenters can test psychological theories of human cognition, behavior, and emotion by systematically creating story microworlds, controlling several variables, and observing the comprehenders' responses. Fourth, discourse contains multiple levels of representation, such as phonemes, words, syntax, propositions, and global messages. The fact that comprehenders are

able to coordinate these multiple levels very quickly is a major achievement that is worthy of scrutiny. And fifth, discourse processing is intertwined with virtually all cognitive functions and processes, including memory, perception, problem solving, and reasoning. When a person is asked to solve a problem, for example, the problem is frequently presented as discourse, and a successful solution presupposes an adequate comprehension of the problem.

The field of discourse processing has grown tremendously during the past decade, which has resulted in several new journals, societies, and conferences. It is beyond the scope of this chapter to provide a comprehensive coverage of the exciting new empirical findings and theoretical developments. We focus here on the comprehension of written text. The scenario to imagine is a college student reading a literary short story for enjoyment, or studying a technical text for an examination. Thus, the emphasis is on written text rather than oral conversation, and on comprehension rather than the production of discourse. We also focus on the representation of meaning, which includes semantics, pragmatics, and the body of knowledge conveyed in the text. The more shallow levels of code (such as phonology, intonation, syntax, and the lexicon of word meanings) are addressed only to the extent that they help clarify how meaning representations are constructed. We recommend Gernsbacher's (1994) *Handbook of Psycholinguistics* for readers who desire a comprehensive coverage of psycholinguistics and all levels of discourse processing. An excellent coverage of oral discourse is provided in books by Clark (1993), Levelt (1989), and Rubin (1995).

Discourse psychologists have investigated a broad array of written texts. At one extreme, researchers investigate naturalistic texts that are written by professional writers for the general public (van Oostendorp & Zwaan 1994). In the narrative genre, the texts have ranged from simple well-formed folktales to literary short stories (Dixon et al 1993, Gerrig 1993, Kreuz & MacNealy 1996, Miall & Kuiken 1994). Texts in the expository genre have frequently covered topics in history (Perfetti et al 1995, Voss & Silfies 1996) and science (Chi et al 1994, Kintsch 1994). These investigations of naturalistic text uncover a representative set of discourse features, patterns, devices, meanings, and comprehension processes that are prevalent in a culture. However, the advantage of ecological validity comes at the cost of losing precise control over the texts' stimulus properties. Consequently, at the other extreme, experimenters carefully craft texts to manipulate independent variables, control for extraneous variables, and satisfy counterbalancing constraints. We call these experimenter-generated materials "textoids" because they are not naturalistic discourse segments that are written to convey an informative or interesting message to a comprehender. Indeed, the texts in far too many experiments are meandering, choppy, pointless, and uninteresting; such texts may impose con-

trol over shallow levels of code (such as word frequency, word meanings, and syntactic composition) but fail miserably in providing control over global coherence and information value. There is the risk that the study of textoids unveils unnatural representations and processing strategies. Discourse psychologists are on solid footing when a hypothesis is confirmed in a sample of naturalistic texts in addition to properly controlled textoids.

The methods of investigating text comprehension are quite diverse (Haberlandt 1994). Sometimes the objective is to study the meaning representations that are established after comprehension is completed. Claims about these mental representations are tested by collecting recall protocols, summary protocols, answers to questions, and various judgments on test statements (e.g. new/old recognition judgments, true/false verification judgments, importance ratings). However, these "off-line" measures are not well suited to capturing the processes and representations that are constructed "on-line" during comprehension. What measures and tasks uncover on-line comprehension processes? This question has been debated at length and is far from settled. One straightforward approach is to collect reading times as readers normally read the text. In eye tracking experiments, the researcher records gaze durations on individual words and patterns of eye movements across the words (Garrod et al 1994, Just & Carpenter 1992, Rayner et al 1994). Self-paced reading times are collected by having the reader press a response key after reading individual text segments, such as words, clauses, sentences, or paragraphs (Haberlandt & Graesser 1985). Although these reading times are natural, the times can sometimes be ambiguous with respect to the contents and types of processes they index. Additional clarity is provided in tasks that periodically interrupt the reader and collect data during comprehension. For example, in a "think aloud" task, the reader expresses ideas that come to mind as each clause in the text is comprehended. The content extracted from think aloud protocols is a very rich source of data for discovering possible comprehension strategies and for testing detailed claims about the representations that enter the reader's consciousness. Researchers have also demonstrated that think aloud protocols are somewhat valid reflections of normal comprehension activities (Chi et al 1994, Ericsson & Simon 1993, Trabasso & Magliano 1996, Zwaan & Brown 1996). However, the protocols do not reliably tap unconscious comprehension processes. Both conscious and unconscious comprehension processes can be tapped in a word-naming task in which readers are periodically interrupted during comprehension and asked to name a test word as quickly as possible. The word-naming latency should be quick if the features of the word closely match a representation that is active in the reader's mind. As an alternative to the word-naming latencies, researchers frequently collect lexical decision latencies on test strings (i.e. whether a sequence of letters forms a word or a

nonword), or word recognition latencies (i.e. whether a test word appeared earlier in the text). Unfortunately, there is a drawback to these tasks that interrupt the reader for data collection: The reader might suffer from "comprehension interruptus" and resort to constructing an unnatural, choppy, shallow representation. Therefore, the rigorous discourse psychologist insists on converging evidence from multiple methods before accepting an empirical claim as valid. Some researchers have advocated a "three-pronged method" that coordinates (*a*) predictions generated by theories, models, and hypotheses; (*b*) data from think aloud protocols; and (*c*) behavioral measures that assess processing time (Graesser et al 1994, Magliano & Graesser 1991, Millis & Graesser 1994, Suh & Trabasso 1993, Zwaan & Brown 1996).

MULTIPLE LEVELS OF DISCOURSE REPRESENTATION

Several levels of discourse representation have been identified by scholars in text linguistics, computational linguistics, sociolinguistics, and literary studies. However, some of these levels have not been embraced by discourse psychologists because they are esoteric or are applicable to a very narrow set of discourse contexts. Most discourse psychologists adopt van Dijk & Kintsch's (1983) distinctions among the *surface code*, the *textbase*, and the referential *situation model*. The surface code preserves the exact wording and syntax of clauses. Comprehenders normally retain the surface code of only the most recent clause unless aspects of this surface code have important repercussions on meaning. The textbase contains explicit text propositions in a stripped-down form that preserves meaning, but not the exact wording and syntax. The textbase also includes a small number of inferences that are needed to establish local text coherence. The situation model is the content or the microworld that the text is about. The situation model for a story refers to the people, spatial setting, actions, and events in the mental microworld. This microworld is constructed inferentially through interactions between the explicit text and background world knowledge.

In addition to these three levels of representation, psychologists normally acknowledge representations and processes at two other levels, which we call the *communication* level and the *text genre* level. The communication level refers to the pragmatic communicative context within which the text is embedded. Thus, the writer prepares the text to communicate ideas to readers (Nystrand 1986), and story narrators communicate episodes to narratees. Regarding text genre, discourse analysts have identified many categories and subcategories of text genre (Biber 1988), such as narration, exposition, description, persuasion, jokes, and so on. A newspaper article, for example, involves quite different structural components, features, and pragmatic ground rules than a joke. All five of these levels contribute to the meaning repre-

sentations that readers build during comprehension. Moreover, it is a profound understatement to say that these various levels interact with one another in complex ways that are not well understood.

To illustrate the five levels of representation, consider the excerpt below that was extracted from the novel *Einstein's Dreams* by Alan Lightman (1993, p. 102):

> A mushy, brown peach is lifted from the garbage and placed on the table to pinken. It pinkens, it turns hard, it is carried in a shopping sack to the grocer's, put on a shelf, removed and crated, returned to the tree with pink blossoms. In this world, time flows backward.

The *text genre* is literary narrative. The excerpt is extracted from the beginning of a chapter, somewhere in the middle of the book. The novel has a series of chapters that describe different fictitious villages in Switzerland in 1905. Each village directly challenges our normal concept of time by transforming a basic assumption in our TIME schema. For example, the citizens in one village know about the future but not the past, which is opposite to one assumption in our TIME schema. In the village described above, time flows backward, which clearly violates the normal forward flow of time, from past to present to future. At the pragmatic *communication* level, the writer or narrator is attempting to unveil fresh insights about time, reality, and life to the reader by violating the normal assumptions about time. The writer has used a well-known literary device called defamiliarization (Miall & Kuiken 1994). That is, prototypical concepts are transformed in an unfamiliar way by stylistic devices, which forces the reader to reinterpret referents and view them in a new perspective. The events in the first two sentences are very difficult to comprehend as they are being read on-line because there are no obvious causal connections between successive events. The sequence of events in this *situation model* is incoherent. Then the third sentence reveals that time flows backward; consequently, the order of events in the explicit text is opposite to the normal flow of events in a generic FRUIT DISTRIBUTION schema. A diligent reader would have to reinterpret the situation model that was constructed from the first two sentences. It is uncertain at this point exactly what deep messages the author wants to communicate by crafting a text with discrepancies among (*a*) the presentation order of events in the text, (*b*) the order of events in a generic FRUIT DISTRIBUTION schema, and (*c*) the chronological order of events in the situation model for that village in Switzerland.

The *textbase* is normally represented as a structured set of *propositions.* A proposition refers to a state, event, or action and may have a truth value with respect to a real or imaginary world. Each proposition contains a *predicate* (e.g. main verb, adjective, connective) and one or more *arguments* (e.g. nouns, embedded propositions). Each argument has a functional role, such as agent,

patient, object, or location. The textbase of propositions is presented below for the first sentence in the example excerpt.

PROP 1: lift (AGENT = X, OBJECT = peach, SOURCE = from garbage)
PROP 2: brown (OBJECT = peach)
PROP 3: mushy (OBJECT = peach)
PROP 4: place (AGENT = X, OBJECT = peach, LOCATION = on table)
PROP 5: pinken (OBJECT = peach)
PROP 6: [in order] to (PROP 4, PROP 5)
PROP 7: and (PROP 1, PROP 4)

The seven propositions have predicates that are verbs (lift, place, pinken), adjectives (brown, mushy), and connectives (in order to, and). The arguments include objects (peach, garbage, table), an unidentified agent (X), and embedded propositions (e.g. PROP 4 and PROP 5 are embedded in PROP 6). Note that the propositional textbase does not capture several features of the *surface code*, such as tense, aspect, voice, and the determinacy of the nouns. For example, the textbase does not capture the fact that the sentence syntax is in the passive voice rather than the active voice. It does not indicate that peach has an indeterminate referring expression (i.e. a peach) whereas table is determinate (i.e. the table).

Separation and Interaction of Levels

Most researchers believe that the five levels of representation exist and are sufficiently distinct for researchers to isolate. However, these beliefs have been challenged. For example, there is not a perfect consensus that there is a separate textbase. Instead, the syntactic composition and lexical items may directly serve as cues or processing instructions on how to construct the situation model, without there being any intermediate textbase of propositions (Gernsbacher 1990, Givón 1992, Perfetti & Britt 1995). Similarly, the reader of a novel may not construct an invisible, virtual writer or storyteller that communicates with the reader, unless there are explicit features in the text that signal that communication level. Instead, the reader may merely become absorbed in the microworld as a voyeur or side participant (Duchan et al 1995, Gerrig 1993). A persistent challenge has been to devise experimental tasks that isolate the separate levels of representation.

Discourse psychologists have collected sentence recognition judgments in an effort to tease apart the surface code, the textbase, and the situation model (Kintsch et al 1990, Schmalhofer & Glavanov 1986, Zwaan 1994). After reading a text, the participants are given a recognition test on the following classes of test sentences: (*a*) the original sentence verbatim, (*b*) a paraphrase of the original sentence, (*c*) a plausible inference with respect to the situation model, and (*d*) a false statement. A subtraction procedure is used to define the surface code (*a* minus *b*), the textbase (*b* minus *c*), and the situation model (*c*

minus *d*). This approach to measuring the three discourse levels has produced theoretically sensible results. For example, there was a rapid decay of the surface code as a function of retention interval and a very slow decay of the situation model, with the textbase in between. When readers believe they are reading literature, the surface code is enhanced, and the situation model is reduced compared with when readers believe they are reading newspaper articles (Zwaan 1994). Therefore, readers are concerned about what is true about the world when they read newspaper articles, whereas they attend to more of the wording and stylistic devices when they read literature. Results such as these suggest that there are natural demarcations among the surface code, the textbase, and the situation model.

Kintsch and his associates have also explored individual differences among readers in an effort to segregate differences between the textbase and the situation model (Kintsch 1994, Mannes 1994, McNamara et al 1995). In McNamara et al, a technical text on the functioning of the heart was studied by students who varied in their background knowledge about the heart (low versus high knowledge). The coherence of the textbase was manipulated by having different versions of the text. Text coherence was enhanced by linking clauses with appropriate connectives and/or by inserting topic sentences, headings, and subheadings at appropriate locations. After studying the texts, the students were tested with tasks that tap the textbase (such as recall for the text) and tasks that tap the situation model (such as difficult questions that require reasoning and problem solving). The results for the low-knowledge readers were compatible with virtually all theories of comprehension. That is, a coherent textbase enhanced performance on measures of both the textbase and the situation model. For high-knowledge readers, however, the pattern of results was more interesting. A coherent textbase slightly enhanced recall but actually lowered performance on tasks that tap the situation model. This cross-over interaction supports the claim that the textbase can be separated from the situation model. Moreover, these results have intriguing implications for education and the writing of textbooks. A coherent textbook improves learning for readers with low knowledge, no matter how the learning is measured. However, readers with an adequate background knowledge may actually benefit from a text with coherence gaps and other obstacles that prevent superficial processing. A coherent text that explicitly lays out the material may give readers with comparatively high knowledge an illusory feeling that they have understood all of the explicit text and its implications, when in fact their representations are imperfect at the deeper situation model.

There have been some lively debates about the interaction of and the time course of constructing the discourse levels. One debate addresses whether the processing of the surface code (which is known to be very quick) is initially

influenced by the other four levels of discourse. According to modularity theory (Fodor 1983), there is an autonomous module for processing syntax, and this module is executed much more quickly than the other discourse levels. The other discourse levels may subsequently override the initial product of the syntax module, but it is syntax that reigns supreme early in the processing stream. According to an interactive theory (Just & Carpenter 1992, MacDonald et al 1994, McClelland & Rumelhart 1986, Whitney et al 1995), the semantic and discourse context can exert its influence early in syntactic parsing. There appears to be some support for modularity theory in analyses of eye tracking data and other measures of on-line processing (Rayner et al 1992), but occasionally a highly constraining semantic or discourse context can have an early influence on parsing (MacDonald 1994, Perfetti & Britt 1995). A similar debate has been pitched at lexical processing. According to modularity theory, the different senses of a word are quickly activated autonomously. In the first sentence of the example text by Lightman (above), the word *table* has at least two senses (e.g. furniture versus organized information on a page), but only the furniture sense is compatible with the situation model. Are both senses automatically activated, or does discourse somehow alter the activation of the two senses? Early research supported modularity theory, but more recent studies have demonstrated that discourse does quickly affect word sense activation (Hess et al 1995, Morris 1994, Rayner et al 1994).

Another debate is pitched at deeper levels of discourse analysis. Sometimes there is a discrepancy between the literal meaning of a sentence (which corresponds to the textbase level) and the meaning that the writer intends to convey (which corresponds to the communication level). A discrepancy between the literal and intended meanings occurs in the case of metaphor (e.g. "All jobs are jails"), irony (e.g. "What lovely weather!" being expressed in a rainstorm), and indirect requests (e.g. "Could you pass the salt?"). It technically is not true that all jobs are jails, so the reader infers that the statement must be metaphorical and that the writer is making some illuminating point. The exclamation "What lovely weather!" directly contradicts the rainy state of the world, so the comprehender infers that the statement is ironical or sarcastic. "Could you pass the salt?" would be an insincere question if the addressee were perfectly capable of passing the salt, so the addressee infers that the speech act is intended as an indirect request for the addressee to perform an action. Early research suggested that there is a two-stage model in which (*a*) the literal meaning was constructed before the intended meaning and (*b*) the intended meaning was constructed only if the literal meaning was implausible in the discourse context (Clark & Lucy 1975). However, subsequent research revealed that intended meanings can be constructed as quickly as literal meanings and do not depend on an implausible literal meaning (Gibbs 1994, Glucksberg et al 1982). In fact,

Gibbs has directly challenged the concept of literal meaning and the claim that a literal meaning is constructed in a discourse context.

Pragmatics and Agents of Communication

It is sometimes claimed that multiple agents, dialogues, and channels of communication are implicitly constructed when texts are comprehended (Bakhtin 1981, Chafe 1994, Clark 1993, Graesser et al 1996, Moffett & McElheny 1995). The agents are capable of speaking, perceiving, knowing, wanting, acting, and experiencing emotions. *Character* agents comprise one ensemble of agents in novels and short stories. These character agents communicate with each other in stories through direct speech (e.g. June told Henry, "I'm pregnant") and indirect speech (June told Henry she was pregnant). *Pragmatic* agents participate in the communicative exchange between the narrator and narratee, or between the writer and reader. It is possible to amalgamate the character agents with the pragmatic agents. In first-person narration, the narrator agent is amalgamated with a character agent (e.g. I woke up one morning and discovered I was pregnant). In second-person narration, there may be an amalgamation of the narrator, the narratee, the reader, and a character in an effort to engage the reader (e.g. You wake up one morning and discover you are pregnant). In third-person narration, there may be a detached, omniscient, all-knowing agent that oversees the storyworld and reports it to the reader (e.g. She woke up one morning and discovered that she was pregnant). The omniscient third-person narrator is invisible to most readers who are not trained in literary studies (Duchan et al 1995, Graesser et al 1996). Discourse psychologists have recently explored how much comprehenders keep track of the knowledge and points of view of the various agents in these multiagent systems (Duchan et al 1995; Graesser et al 1996; Keysar 1994, 1996; Schober 1995, Stein & Liwag 1996). For example, Graesser et al (1996) reported that college students are quite good at keeping track of *who said what* and *who knows what* in literary short stories, except in the case of the third-person narrators. Keysar (1996) reported that readers are better able to keep track of the intentions and knowledge of the speakers than the addressees in embedded dialogues.

Pragmatic principles facilitate communication between agents when messages are composed (Clark 1993, Givón 1992, Grice 1975). Agents that both send and receive messages must mutually agree that these ground rules are operating. When the principles are violated, comprehension time increases or misunderstandings occur.

1. *Monitor common ground and mutual knowledge.* The writer should keep track of words, ideas, and entities that the reader already knows. If some-

thing new is being introduced, it should be signaled syntactically and embellished with adjectives, phrases, or examples.

2. *Use discourse cues to distinguish "given" versus "new" information.* For example, the given information is typically included in the subject noun-phrase of a clause and the first clause of multiclause sentences, whereas the new information is in the verb-phrase and additional clauses.
3. *Use discourse cues to signal important information.* For example, the first sentence in a paragraph should convey a main point and serve as an umbrella for subsequent sentences in the paragraph.
4. *Make true claims about the situation model under consideration.* In expository text, claims should be true about the world in general. In narrative fiction, the claims should be true about the fictitious microworld.
5. *The incoming sentence should be relevant to the previous discourse context.* New topics, subtopics, and episodes need to be flagged with discourse cues, such as subtitles and transitional phrases (e.g. Another point is that…, The next morning…).
6. *The order of mentioning events should correspond to the chronological order of events in the situation model.* This principle is violated in the first two sentences of the Lightman excerpt, but then the third sentence explicitly declares that time flows backwards.
7. *Statements should not contradict one another.*

These pragmatic principles are automatized and unconscious in the minds of most readers, at least those who do not work in a communication profession. Indeed, the principles are so entrenched that some readers never regard it as an option that a writer would express ideas that are false, contradictory, or irrelevant; they faithfully accept pretty much whatever the writer expresses. Beck et al (1996) has implemented a year-long program in the classroom (called Questioning the Author) that trains students to question the rationale and evidence behind particular statements expressed by authors. The students imagine a real flesh-and-blood writer and ask questions such as "What is the author trying to say?" and "What did the author mean by that?" This inquisitive strategy produces a more elaborate representation at the communication and situation model levels. Without this mindset, most readers assume that the writer is faithfully following the pragmatic principles.

Studies have shown that it is difficult for readers to detect anomalous statements (i.e. those that are false, irrelevant, or contradictory) in expository texts on unfamiliar topics (Graesser & McMahen 1993, Otero & Kintsch 1992). Readers miss these anomalies and assume that they understand the text in the absence of such problematic textual features (Glenberg & Epstein 1987). It apparently takes a large amount of background knowledge for a reader to detect anomalous information. However, anomalies are well remembered

when they are detected (Albrecht & O'Brien 1993, Davidson 1994, Graesser et al 1979).

PSYCHOLOGICAL MECHANISMS IN THEORIES OF COMPREHENSION

Psychological models of discourse processing have specified in rich detail how the multilevel meaning representations are built during comprehension. Discourse psychologists have consistently recognized the need to ground these complex models in general theories of cognition. This section briefly enumerates the key cognitive components, processes, and factors that have frequently been adopted by discourse psychologists.

Cognitive Components

1. *Knowledge structures.* The knowledge in texts and in packages of world knowledge are represented as a network of *nodes* (i.e. concepts, referents, propositions) that are interconnected by relational *arcs* (Graesser & Clark 1985, van Dijk & Kintsch 1983). One source of comprehension difficulty lies in the amount of background knowledge of the reader.
2. *Spreading activation of nodes in knowledge networks.* When a node in a network is activated, activation spreads to neighboring nodes, then neighbors of neighbors, and so on. The activation level of a node decreases as a function of the number of arcs between the originally activated node and another node in the network (Anderson 1983).
3. *Memory stores.* There are three memory stores in most discourse models: short-term memory (STM), working memory (WM), and long-term memory (LTM). As a gross approximation, STM holds the most recent clause being comprehended and WM holds about two sentences. Important information is actively recycled in WM (Fletcher & Bloom 1988, Kintsch & van Dijk 1978, Trabasso & Magliano 1996).
4. *Discourse focus.* Consciousness and focal attention is concentrated on one or two nodes in the discourse representation (Chafe 1994, Givón 1992, Grosz & Sidner 1986, Sanford & Garrod 1981). In the situation model for a narrative text, the discourse focus is analogous to a mental camera that zooms in on particular characters, objects, actions, events, and spatial regions (Bower 1989).
5. *Resonance.* The content (i.e. cues, features, nodes) that resides in the discourse focus, STM, and WM may match highly with content that was presented earlier in the text or with other content in LTM. If so, there is resonance with the content in LTM, and the information in LTM gets activated (Albrecht & O'Brien 1993, McKoon & Ratcliff 1992, McKoon et al 1996, Myers et al 1994, O'Brien et al 1995). The content of WM on p.

124 in a novel could quickly activate the content on p. 14 via resonance, without activating any of the content between pp. 15 and 123.

6. *Activation, inhibition, and suppression of nodes.* As sentences are comprehended, nodes in the discourse structure and LTM are activated, strengthened, inhibited, and suppressed (Gernsbacher 1990, Kintsch 1988). The primary goal of some discourse models is to explain the fluctuations in activation values of discourse nodes during the dynamic processes of comprehension.
7. *Convergence and constraint satisfaction.* Discourse nodes receive more strength of encoding to the extent that they are activated by several information sources and to the extent that they mesh with the constraints of other information sources (Graesser & Clark 1985, Kintsch 1988, MacDonald et al 1994).
8. *Repetition and automaticity.* Repetition increases the speed of accessing a knowledge structure and the nodes within the structure. Thus, familiar words are processed faster than unfamiliar words. The nodes in an automatized package of world knowledge are holistically accessed and used at little cost to the resources in WM.
9. *Explanations.* Memory for information is enhanced when the reader constructs causal explanations of why events in the situation model occur and why the writer expresses information (Chi et al 1994, Graesser et al 1994, Pressley et al 1988, Trabasso & Magliano 1996, Zwaan & Brown 1996). Readers actively seek these explanations during reading (Graesser at al 1994).
10. *Reader goals.* The goals of the reader influence text comprehension and memory (Graesser et al 1994, Lorch et al 1995, Zwaan et al 1995b). Reading a novel for enjoyment is rather different from reading it to take a university exam.

Cognitive Models and Architectures

Discourse psychologists have developed some sophisticated quantitative and computational models of text comprehension (Britton & Graesser 1996, Golden & Rumelhart 1993, Just & Carpenter 1992, Kintsch 1988, St. John 1992). These models specify the representations, processes, and interactive mechanisms in sufficient detail to simulate complex patterns of comprehension data. The most fine-grained models simulate the creation, activation, inhibition, and suppression of each node in the discourse representation, as text is dynamically comprehended, word-by-word or clause-by-clause. Consider the activation strength of a word node at a particular point in a text. That strength value should predict latencies in such tasks as word naming, word recognition, and lexical decision. Word reading times and the gaze durations on words should correlate with the number of processing cycles that the model takes to interpret

the word. Memory for text propositions should correlate with the cumulative strength of activation for the proposition across the entire text. In fact, these complex models have had some success in simulating such data.

Two cognitive models have dominated most of the efforts in simulating discourse data: the CAPS/READER model developed by Just & Carpenter (1992) and the construction-integration model developed by Kintsch (1988). Goldman et al (1996) developed a hybrid model that combines these two models. The CAPS/READER model adopts a production system architecture (Anderson 1983) for creating, updating, and removing nodes in WM and LTM. The production system contains a set of production rules with an "IF<condition C>,THEN<action A>" format; if the content of WM matches condition C, then the cognitive or physical action A is performed. Condition C may consist of an arbitrarily complex set of substates. There also is a threshold criterion for a condition, such that the condition is satisfied if the aggregate activation value of all its substates meets the threshold. Therefore, the production rules in CAPS are hardly brittle, discrete, and simple. The set of production rules are evaluated in parallel in each cycle of processing. Those rules that meet the threshold of activation end up executing various actions, such as scanning explicit input, modifying activation values of nodes in WM, changing the load on WM, strengthening nodes in LTM, and producing output.

Kintsch's (1988) construction-integration model adopts a connectionist (neural network) architecture of cognition (McClelland & Rumelhart 1986). During the construction phase, an incoming clause very quickly adds to WM a set of nodes that corresponds to words, referents, textbase propositions, and the situation model. These new nodes are combined with the previous content of WM. Suppose, for illustration, that WM has a total of N nodes when the incoming clause is comprehended. In the spirit of connectionism, there is a set of weights [$N \times (N - 1)$] that designates how much each node (M) would activate or inhibit each of the other nodes if node M were in fact activated. There is a separate weight space for the surface code, the textbase, and the situation model, along with weights that connect nodes between levels. Whereas the construction phase is accomplished quickly and automatically, the integration phase is more time consuming and spans several processing cycles. The integration phase begins as soon as one or more of the N nodes becomes activated. An activated node spreads activation or inhibition to the other nodes according to the weights in the weight space. The spreading activation continues over several processing cycles until the connectionist network settles on a stable set of activation values for the entire set of discourse nodes. The activation strength of a particular node is modified dynamically over time and over sentences in the text.

REFERRING EXPRESSIONS

Referring expressions are nouns, pronouns, and noun-phrases that refer to an entity or proposition in the textbase, situation model, or world. An *anaphoric* referring expression refers to a node that was mentioned previously in the text, whereas *cataphoric* expressions refer to future text nodes, and *deictic* expressions point to the world. Anaphoric expressions have received the most attention by discourse psychologists (Garrod et al 1994, Gernsbacher 1990, Greene et al 1992, Marslen-Wilson et al 1993, McKoon et al 1996, Sanford & Garrod 1981), but recently some serious attention has turned to cataphoric expressions (Gernsbacher & Jescheniak 1995) and deictic expressions (Duchan et al 1995, Mauner et al 1995). A concrete example of anaphora is when a reader encounters the pronoun *he* in the middle of a novel. Who does *he* refer to? The process of resolving the referent for *he* consults all levels of discourse.

The selection of referring expressions normally conforms to a small set of simple rules (Chafe 1994, Gernsbacher 1990, Givón 1992, Sanford & Garrod 1981). When a new entity is first introduced in a text, the referring expression contains (*a*) the indefinite determiner *a* or *an,* (*b*) a richly specified noun, and (*c*) a descriptive set of adjectives and prepositional phrases (e.g. *A big bad wolf with brown fur*). A pronoun is appropriate when the entity has already been mentioned in the text and is also in the discourse focus. When the entity has already been introduced, but is not in the discourse focus, an appropriate referring expression contains a definite determiner (e.g. *the, that, this*) and the noun (*the wolf*); the noun is sometimes at a more abstract level of specification (*the animal* instead of *the wolf*). If two discourse entities are similar, the referring expression needs to be sufficiently rich to distinguish them. Thus, fewer words and less specificity are needed to the extent that the entity is in the reader's working memory, and fewer yet when the entity is in the reader's focal attention. More explicit cues are needed in referring expressions when information needs to be created from scratch, to be dredged from LTM, or to distinguish entities. When these simple rules are violated, it takes longer to compute the referent of a referring expression, and comprehension may break down (Gernsbacher 1990, Gordon & Chan 1995).

Gernsbacher's (1990) *structure building framework* accounts for much of the experimental data on the processing of anaphoric references. Three subprocesses occur when discourse representations are constructed on-line. Readers first "lay a foundation" by constructing memory nodes when a new topic is introduced. The foundation gets elaborated by "mapping on information" from subsequent text that is relevant to the topic. However, the incoming text may not be relevant, so the reader shifts attention to "initiate a new substructure" or to lay an entirely new foundation. The first two subprocesses explain why the discourse focus is on the first character in sentences such as *Tina beat Lisa in*

the state tennis match. Tina is the discourse focus, so *Tina* should have a higher level of activation and be accessed faster than *Lisa;* if the next word were *she,* the reader would bind the pronoun to *Tina* rather than *Lisa.* Gernsbacher's experiments have supported this prediction, dubbed the "advantage of first mention." Her experiments also support a phenomenon called "the advantage of clause recency," which predicts that information from the most recent clause in a sentence is more accessible than information from an earlier clause. The word *oil* is more accessible immediately after comprehending the sentence *Now that artists are working fewer hours,* oil *prints are rare* than the sentence *Now that artists are working in* oil, *prints are rare.* Comprehenders represent each clause in these two-clause sentences in its own mental substructure. When a new clause arrives, the old substructure from clause 1 is abandoned and attention shifts to the new substructure associated with clause 2.

Our discussion of anaphora presents a simple sketch of the processing and appropriate composition of the referring expressions. However, there are times when matters are far more complex (Clark 1993, Garrod et al 1994, Greene et al 1992, Marslen-Wilson et al 1993, McKoon et al 1996). In some cases, for example, readers do not bother fetching a referent for an anaphor because the referent is vague, indeterminate, difficult to compute, or nonexistent (such as the pronoun *it* in technical documents).

CONNECTING STATEMENTS IN DISCOURSE

The explicit statements in a text need to be connected conceptually if the text is to be regarded as coherent. Local coherence is achieved if the reader can connect the incoming statement to information in the previous sentence or WM. Global coherence is achieved if the incoming statement can be connected to a text macrostructure or to information much earlier in the text that no longer resides in WM. Readers normally attempt to achieve coherence at both the local and global levels (Albrecht & O'Brien 1993, Graesser et al 1994, Hakala & O'Brien 1995, Hess et al 1995, Long et al 1996, Myers et al 1994, O'Brien & Albrecht 1992, Sanford & Garrod 1981, Singer et al 1994, Trabasso & Magliano 1996). Suppose, for example, that a character is described as a vegetarian early in the text and that much later the text states that the character ate a hamburger. The contradiction could only be detected if the reader were attempting to achieve global coherence. Reading times have been found to increase for such contradictory statements under conditions in which the statement *X is a vegetarian* has no local connections to *X ate a hamburger* (Albrecht & O'Brien 1993, Hakala & O'Brien 1995, Myers et al 1994). This increase in reading time would not occur if text comprehension was driven entirely by local connections. However, it is important to acknowledge that attempts at achieving global coherence will diminish and local coherence will

dominate if the text is incoherent, if the reader is unmotivated, or if the reader has a low WM span (Graesser et al 1994, Hess et al 1995, Whitney et al 1991). Comprehension suffers substantially when neither local nor global coherence can be achieved, as in the case of the first two sentences of the Lightman excerpt. As a general underlying principle, readers attempt to achieve the most global level of understanding that can be achieved given the text composition, the reader's knowledge base, and the reading goals.

Several dimensions of conceptual continuity link an incoming statement to the previous discourse context (Chafe 1994, Graesser & Clark 1985, Grimes 1975, Halliday & Hasan 1976, Mann & Thompson 1986, Sanders et al 1992). The reading time for a sentence increases when there are breaks in continuity on one or more of these dimensions. One dimension of continuity is *argument overlap* (Kintsch & van Dijk 1978). Continuity is achieved if there is a noun-phrase argument in the incoming statement that overlaps an argument within any textbase proposition in WM. In the Lightman example, the pronoun *it* in sentence 2 refers to the argument *peach* in sentence 1; therefore the two sentences are connected by argument overlap. Discourse psychologists have sometimes regarded argument overlap as the primary dimension for establishing text coherence (Kintsch & van Dijk 1978, McKoon & Ratcliff 1992). The reading time for a sentence in text does normally increase if it fails to share an argument with any textbase proposition in WM (Kintsch & van Dijk 1978). However, argument overlap is sometimes not a major dimension if the text is in the narrative genre or is fortified by rich background knowledge (Zwaan et al 1995a,b).

Zwaan proposed an *event indexing model* to account for the reader's construction of a multithreaded situation model while reading simple stories and literary short stories (Zwaan et al 1995a,b). According to this model, the reader monitors five conceptual dimensions during reading: the protagonist, temporality, spatiality, causality, and intentionality (i.e. character goals). A break in continuity may occur on any of these dimensions when an incoming statement is read.

1. *Spatial discontinuity.* The incoming event occurs in a spatial setting that is different from the prior event. Sometimes this is manifested by a transitional phrase (e.g. *Back at the ranch,…*).
2. *Temporal discontinuity.* The incoming event occurs much later in time (e.g. *The next day.…*) or is part of a flashback.
3. *Causal discontinuity.* The incoming event is not causally related to the prior text (as in the Lightman example).
4. *Intentional discontinuity.* The incoming event is embedded in a character's plan that is different from the local discourse context.

5. *Protagonist discontinuity*. The incoming event has a character that is different from the characters in the previous event.

An incoming event in the story may have discontinuities on more than one of these five dimensions. Zwaan et al (1995a) reported that the reading time for an explicit event in a literary story increased as a function of the number of dimensions with discontinuities and that each dimension had its own unique impact on reading time. The event indexing model is also compatible with a large body of research that has examined each of these dimensions one at a time. That is, researchers have confirmed that readers construct a situation model that monitors spatiality (Glenberg et al 1987, Haenggi et al 1995, Morrow et al 1987, O'Brien & Albrecht 1992, Rinck et al 1996), temporality (Bestgen & Vonk 1995, Ohtsuka & Brewer 1992, Zwaan 1996), causality (Fletcher & Bloom 1988, Keenan et al 1984, Klin & Myers 1993, Magliano et al 1993, Millis & Graesser 1994, Singer et al 1992, Trabasso & van den Broek 1985, Trabasso & Magliano 1996, van den Broek & Lorch 1993), intentionality (Dopkins 1996, Long et al 1992, Suh & Trabasso 1993), and properties of the protagonist (Albrecht & O'Brien 1993, Hakala & O'Brien 1995, Myers et al 1994).

The surface code delivers critical cues to the reader, who is actively monitoring the dimensions of conceptual continuity. As suggested above, transitional phrases are quite important signals for spatial and temporal discontinuity (e.g. *Back at the ranch,… The next morning.*). Verb tense and aspect are also important cues for situating events on a chronological time line. One very important class of cues for the dimensions of temporality, causality, and intentionality is that of connectives, such as *before, after, during, and, then, because, in order to,* and *so that.* Discourse researchers have investigated reading times and recall for successive text statements that are explicitly linked by causal connectives (such as *because*) and temporal connectives (such as *and*), versus no connective (Caron et al 1988, Deaton & Gernsbacher 1996, Millis et al 1993, Millis & Just 1994, Murray 1995). Compared with temporal connectives and no connective, the causal connectives facilitate reading time and later recall for the statements. However, these effects apparently do not always occur. Memory is not facilitated if the events linked by a causal connective already have a very strong causal relationship or no causal relationship. Thus, a writer cannot simply slap in a causal connective and expect it to work its magic. The causal connective needs to mesh with the pragmatic levels of the discourse context in an incisive way. For example, according to Ford's (1994) analysis of causal connectives in speech and writing, the connective *because* (in an expression *X because Y*) is appropriate when (*a*) event *X* violates a norm or deviates from shared expectations and (*b*) *Y* explains the anomaly (e.g. *The boss was absent because he was accused of sexism*).

Discourse psychologists have only a rudimentary understanding of how the various strands of the situation model are constructed during comprehension. Additional research needs to examine interactions among temporality, spatiality, causality, intentionality, and the protagonists. There needs to be a more detailed understanding of how these strands of conceptual continuity are furnished by the surface code, assumptions about communication, and text genre. Hopefully, future researchers will explore these conceptual jungles with the same care and detail with which they have investigated the lexicon, syntax, and semantics during the past three decades.

KNOWLEDGE-BASED INFERENCES

Discourse psychologists have developed and tested models that predict what inferences are generated on-line during comprehension (Graesser & Bower 1990, Graesser et al 1994, McKoon & Ratcliff 1992). When an adult reads a novel, for example, the following classes of knowledge-based inferences are potentially generated: The goals and plans that motivate characters' actions, character traits, characters' knowledge and beliefs, character emotions, causes of events, the consequences of events and actions, properties of objects, spatial contexts, spatial relationships among entities, the global theme or point of the text, the referents of nouns and pronouns, the attitudes of the writer, and the appropriate emotional reaction of the reader. It is conceivable that readers generate all these inferences on-line. In essence, the mind would construct a high-resolution mental videotape of the situation model, along with details about the mental states of characters and the communicative exchange between the writer and reader. However, discourse psychologists are convinced that only a subset of these inferences are generated on-line. Why? Because the generation of all these inferences would create a computational explosion problem, because WM has limited resources, and because reading is accomplished too quickly for some time-consuming inferences to be generated. Which of these classes of inferences are constructed on-line? That is the central question.

Suppose an adult read a simple fairy tale that contained two successive actions in the middle of the text: "The dragon was dragging off the girl. A hero came and fought the dragon." There are five classes of inferences that might be encoded when the second sentence is read:

1. *Superordinate goal (motive).* The hero wanted to rescue the girl.
2. *Subordinate goal or action.* The hero threw a spear.
3. *Causal antecedent.* The girl was frightened.
4. *Causal consequence.* The hero married the girl.
5. *Static property.* The dragon has scales.

Experiments can be designed to assess which of these inferences are encoded. For example, after reading the second sentence, the reader would quickly complete a word-naming task (or alternatively a lexical decision task) and receive a test word extracted from one of these inferences (i.e. rescue, throw, fright, marry, and scales). The same words would also be tested in an unrelated context to obtain a measure of inference encoding from the word-naming latencies: [latency (unrelated context)—latency (inference context)]. The inference encoding score should be above zero to the extent that the inference is generated on-line. In a properly designed study, the words in these inference classes would be equilibrated on several variables, such as number of letters, number of syllables, word frequency, syntactic class, free association norms with words in the text, and the proportion of readers in a normative group who articulate the inference in a think aloud task (Graesser et al 1994, Kintsch 1988, Long et al 1992, Magliano et al 1993, Millis & Graesser 1994). All things being equal, which classes of inference should have the strongest encoding during comprehension?

Existing models make quite different predictions about which of the five classes of inferences are encoded on-line. At one extreme, there is a *promiscuous inference generation* position that predicts that all five classes are encoded. This is a strawperson position, however, for reasons discussed above. At the other extreme, there is a textbase position that predicts that none of the inferences are encoded (Kintsch & van Dijk 1978). Local text coherence could be established at the textbase level by virtue of argument overlap (i.e. *dragon* appears in both sentences) so there would be no need to construct a situation model. According to McKoon & Ratcliff's (1992) *minimalist hypothesis*, the causal antecedents would be the only class of inferences among the five that might be encoded with any consistency. The other four classes are elaborative inferences that are encoded only if the reader has a special comprehension goal that is tuned to such levels. The minimalist hypothesis assumes that the reader generates only those inferences that are needed to establish local text coherence (i.e. either causal antecedents or none at all) and that are readily available in memory (i.e. in WM or highly active in LTM).

Graesser et al (1994) has argued that the available research on inference generation supports a *constructionist theory* rather than the above three positions (as well as others that will not be addressed here). The constructionist theory assumes that readers encode three sets of inferences, namely (*a*) inferences that address the readers' comprehension goals, (*b*) inferences that explain *why* events, actions, and states occur, and (*c*) inferences that establish coherence in the situation model at local and global levels. From the standpoint of the five classes of inferences in the above example, it is the explanation assumption that offers distinctive predictions. The explanation-based in-

ferences include superordinate goals and causal antecedents, but not subordinate goals/actions, causal consequences, and static properties. It is the superordinate goals and the causal antecedents that answer *why* an action or event occurs (Graesser et al 1991). For example, when asked "Why did the hero fight the dragon?" a reasonable answer would include the superordinate goal (in order to rescue the girl) and the causal antecedent (because the girl was frightened), but not the other three classes (i.e. in order to throw the spears, in order to marry the girl, because the dragon has scales). The subordinate goals/actions and static properties are minor ornaments that merely embellish the core plot and explanations of the plot. Regarding the causal consequences, it is too difficult for readers to forecast multiple hypothetical plots with new plans of characters and long event chains into the future. Most of the forecasts that readers generate to "What happens next?" questions end up being wrong (Graesser & Clark 1985), so the readers would be uselessly spinning their wheels if they did generate many causal consequences. According to the constructionist model, the only causal consequences that are generated on-line are the achieved superordinate goals of character actions (i.e. the hero in fact did rescue the girl), emotional reactions of characters to actions and events, and consequences that are highly activated and constrained by prior context (see Keefe & McDaniel 1993, van den Broek 1990). Analyses of think aloud protocols in fact confirm that most readers (and good comprehenders in particular) generate more explanation-based inferences than predictions and associative elaborations in simple stories (Trabasso & Magliano 1996), in literary short stories (Zwaan & Brown 1996), and in technical expository texts (Chi et al 1994). Experiments that have collected word-naming and lexical decision latencies confirm the constructionist theory's predictions that readers generate superordinate goals much more than subordinate goals/actions (Long et al 1992), and causal antecedents much more than causal consequences (Magliano et al 1993, Millis & Graesser 1994, Potts et al 1988).

We suspect that each of the above models is correct in certain conditions. The textbase position and minimalist hypotheses are probably correct when the reader is very quickly reading the text, when the text lacks global coherence, and when the reader has very little background knowledge. The constructionist theory is on the mark when the reader is attempting to comprehend the text for enjoyment or mastery at a more leisurely pace, when the text has global coherence, and when the reader has some background knowledge. The promiscuous inference generation model may even be valid when a literary scholar is savoring a good short story at a very slow cruise.

There are many gaps in our understanding of the generation of knowledge-based inferences. We need to analyze the precise time course of constructing, maintaining, and modifying particular classes of inferences (Keefe &

McDaniel 1993, Kintsch 1988, Magliano et al 1993). Some inferences may slowly emerge as text is received rather than discretely popping in when a particular statement is comprehended. Very little research has examined global inferences, such as themes, points, morals, and attitudes of the writer (Long et al 1996, Seifert et al 1986). Global inferences have tentacles to many elements in the text and span large stretches of text. These global inferences are more difficult to study than inference classes that discretely pop in at a particular locus in the text. There also needs to be much more work on how inference generation is influenced by readers who differ in comprehension skill, working memory span, and other psychological attributes (Dixon et al 1993, Haenggi et al 1995, Perfetti et al 1996, Singer & Ritchot 1996, Whitney et al 1991).

SUMMARY AND CONCLUSIONS

Discourse psychologists have developed sophisticated models of text comprehension during the short 25-year history of the field. These models have specified the cognitive representations and processes that participate in the construction of meaning. Five levels of representation are important in the construction of these meaning representations: the surface code, the propositional textbase, the referential situation model, the communication between writer and reader, and the text genre. Discourse psychologists have explained the processing of these five levels by grounding the research in general cognitive theories. For example, discourse processing theories have postulated the existence of multiple memory stores (STM, WM, LTM), the process of spreading activation in knowledge networks, production systems with if-then rules, connectionist neural networks, and constraint satisfaction. However, the models developed by discourse psychologists furnish distinctive predictions that cannot simply be reduced to other areas of psychology.

This chapter has examined three phenomena that have been extensively investigated by discourse psychologists: The processing of referring expressions, the connection of statements in text, and the encoding of knowledge-based inferences. Readers execute these processes in an effort to achieve coherence at local and global levels and to explain why information is mentioned in the text. Discourse psychologists have tested models of these phenomena by collecting eye-tracking data, self-paced reading times, word-naming latencies on test words that are interspersed in text, recall and recognition memory measures, think aloud protocols, and data from dozens of other experimental tasks. There have been some attempts to assess differences among readers concerning background knowledge about the text topics, working memory span, and general comprehension skill. Some researchers have tested their models on naturalistic texts, whereas others have focused on experi-

menter-generated textoids that impose some control over particular variables. In some cases, complex patterns of data have been simulated by sophisticated mathematical and computational models.

The empirical and theoretical progress in discourse processing has had some straightforward applications for improving reading, education, text design, and social interaction. For example, Britton & Gulgoz (1991) used Kintsch & van Dijk's 1978 model to guide the revision of technical expository texts. The original texts were naturalistic samples of texts that periodically had problematic referring expressions and coherence gaps. The Kintsch & van Dijk model identified points in the text where these problems occurred. The texts were revised by clarifying referents of referring expressions (with definitions and/or examples) and by making explicit some important connections between different parts of the text and some critical bridging inferences in the original text. These theory-based revisions dramatically improved memory for the texts when adults were given a delayed-recall test. Moreover, these theory-guided revisions improved memory much more than revisions by writers for *Time* magazine. It was theory that prevailed in improving the memorability of the texts. This is a heart-warming confirmation of one of our favorite quotes from Kurt Lewin (1951, p. 169): "There is nothing so practical as a good theory."

Literature Cited

Albrecht JE, O'Brien EJ. 1993. Updating a mental model: maintaining both local and global coherence. *J. Exp. Psychol. Learn. Mem. Cogn.* 19:1061–70

Anderson JR. 1983. *The Architecture of Cognition.* Cambridge, MA: Harvard Univ. Press

Bakhtin MM. 1981. *The Dialogic Imagination.* Ed./Transl. C Emerson, M Holquist. Texas: Texas Univ. Press

Beck IL, McKeown MG, Worthy J, Sandora CA, Kucan L. 1996. Questioning the author: a year-long classroom implementation to engage students with text. *Elem. Sch. J.* In press

Bestgen Y, Vonk W. 1995. The role of temporal segmentation markers in discourse processing. *Discourse Process.* 19:385–406

Biber D. 1988. *Variation Across Speech and Writing.* Cambridge: Cambridge Univ. Press

Bower GH. 1989. Mental models in text understanding. In *Cognition in Individual and Social Contexts,* ed. AF Bennett, KM McConkey, pp. 129–44. Amsterdam: Elsevier

Britton BK, Graesser AC, eds. 1996. *Models of Understanding Text.* Mahwah, NJ: Erlbaum

Britton BK, Gulgoz S. 1991. Using Kintsch's computational model to improve instructional text: effects of repairing inference calls on recall and cognitive structures. *J. Educ. Psychol.* 83:329–45

Caron J, Micko HC, Thuring M. 1988. Conjunctions and recall of composite sentences. *J. Mem. Lang.* 27:309–23

Chafe W. 1994. *Discourse, Consciousness, and Time.* Chicago: Univ. Chicago Press

Chi MTH, de Leeuw N, Chiu M, LaVancher C. 1994. Eliciting self-explanations improves understanding. *Cogn. Sci.* 18:439–77

Clark HH. 1993. *Arenas of Language Use.* Chicago: Univ. Chicago Press

Clark HH, Lucy P. 1975. Understanding what is meant from what is said: a study in conversationally conveyed requests. *J. Verbal Learn. Verbal Behav.* 14:56–72

Davidson D. 1994. Recognition and recall of irrelevant and interruptive atypical actions in script-based stories. *J. Mem. Lang.* 33:757–75

Deaton JA, Gernsbacher MA. 1996. Causal conjunctions and implicit causality cue

mapping in sentence comprehension. *J. Mem. Lang.* In press
Dixon P, Bortolussi M, Twilley LC, Leung A. 1993. Literary processing and interpretation: towards empirical foundations. *Poetics* 22:5–34
Dopkins S. 1996. Representation of superordinate goal inferences in memory. *Discourse Process.* 21:85–104
Duchan JF, Bruder GA, Hewitt LE, eds. 1995. *Deixis in Narrative: A Cognitive Science Perspective.* Hillsdale, NJ: Erlbaum
Ericsson KA, Simon HA. 1993. *Protocol Analysis: Verbal Reports as Data.* Cambridge, MA: MIT Press. 2nd ed.
Fletcher CR, Bloom CP. 1988. Causal reasoning in the comprehension of simple narrative texts. *J. Mem. Lang.* 19:70–80
Fodor JD. 1983. *Modularity of Mind.* Cambridge, MA: MIT Press
Ford CE. 1994. Dialogic aspects of talk and writing: *because* on the interactive-edited continuum. *Text* 14:531–54
Garrod S, Freudenthal D, Boyle E. 1994. The role of different types of anaphor in the on-line resolution of sentences in a discourse. *J. Mem. Lang.* 33:39–68
Gernsbacher MA. 1990. *Language Comprehension as Structure Building.* Hillsdale, NJ: Erlbaum
Gernsbacher MA, ed. 1994. *Handbook of Psycholinguistics.* New York: Academic
Gernsbacher MA, Jescheniak JD. 1995. Cataphoric devices in spoken discourse. *Cogn. Psychol.* 29:24–58
Gerrig RJ. 1993. *Experiencing Narrative Worlds.* New Haven, CT: Yale Univ. Press
Gibbs RW. 1994. *The Poetics of Mind.* Cambridge: Cambridge Univ. Press
Givón T. 1992. The grammar of referential coherence as mental processing instructions. *Linguistics* 30:5–55
Glenberg AM, Epstein W. 1987. Inexpert calibration of comprehension. *Mem. Cogn.* 15: 84–93
Glenberg AM, Meyer M, Lindem K. 1987. Mental models contribute to foregrounding during text comprehension. *J. Mem. Lang.* 26:69–83
Glucksberg S, Gildea P, Bookin HB. 1982. On understanding nonliteral speech. Can people ignore metaphor? *J. Verbal Learn. Verbal Behav.* 21:85–98
Golden RM, Rumelhart DE. 1993. A parallel distributed processing model of story comprehension and recall. *Discourse Process.* 16:203–37
Goldman SR, Varma S, Cote N. 1996. Extending capacity-constrained construction integration: toward "smarter" and flexible models of text comprehension. See Britton & Graesser 1996, pp. 73–114
Gordon PC, Chan D. 1995. Pronouns, passives, and discourse coherence. *J. Mem. Lang.* 34:216–31
Graesser AC, Bower GH, eds. 1990. *The Psychology of Learning and Motivation: Inferences and Text Comprehension.* San Diego, CA: Academic
Graesser AC, Bowers CA, Bayen UJ, Hu X. 1996. Who said what? Who knows what? Tracking speakers and knowledge in narrative. In *Narrative Perspective: Cognition and Emotion,* ed. W van Peer, E Andriga, D Schram, E Tan. In press
Graesser AC, Clark LF. 1985. *Structures and Procedures of Implicit Knowledge.* Norwood, NJ: Ablex
Graesser AC, Gordon SE, Sawyer JD. 1979. Memory for typical and atypical actions in scripted activities: test of a script + tag hypothesis. *J. Verbal Learn. Verbal Behav.* 18:319–32
Graesser AC, Lang KL, Roberts RM. 1991. Question answering in the context of stories. *J. Exp. Psychol. Gen.* 120:254–77
Graesser AC, McMahen CL. 1993. Anomalous information triggers questions when adults solve problems and comprehend stories. *J. Educ. Psychol.* 85:136–51
Graesser AC, Singer M, Trabasso T. 1994. Constructing inferences during narrative text comprehension. *Psychol. Rev.* 101: 371–95
Greene SB, McKoon G, Ratcliff R. 1992. Pronoun resolution and discourse models. *J. Exp. Psychol. Learn. Mem. Cogn.* 18: 266–83
Grice HP. 1975. Logic and conversation. In *Syntax and Semantics: Speech Acts,* ed. P Cole, JL Morgan, pp. 3:41–58. San Diego, CA: Academic
Grimes J. 1975. *The Thread of Discourse.* The Hague: Mouton
Grosz BJ, Sidner CL. 1986. Attention, intentions, and the structure of discourse. *Comput. Linguist.* 12:175–204
Haberlandt K. 1994. Methods in reading research. See Gernsbacher 1994, pp. 1–31
Haberlandt K, Graesser AC. 1985. Component processes in text comprehension and some of their interactions. *J. Exp. Psychol. Gen.* 114:357–74
Haenggi D, Kintsch W, Gernsbacher MA. 1995. Spatial situation models and text comprehension. *Discourse Process.* 19: 173–99
Hakala CM, O'Brien EJ. 1995. Strategies for resolving coherence breaks in reading. *Discourse Process.* 20:167–86
Halliday MAK, Hasan R. 1976. *Cohesion in English.* London: Longmans
Hess DJ, Foss DJ, Carroll P. 1995. Effects of global and local context on lexical process-

ing during language comprehension. *J. Exp. Psychol. Gen.* 124:62–82
Just MA, Carpenter PA. 1992. A capacity theory of comprehension: individual differences in working memory. *Psychol. Rev.* 99:122–49
Keefe DE, McDaniel M. 1993. The time course and durability of predictive inferences. *J. Mem. Lang.* 32:446–63
Keenan JM, Baillet SD, Brown P. 1984. The effects of causal cohesion on comprehension and memory. *J. Verbal Learn. Verbal Behav.* 23:115–26
Keysar B. 1994. The illusory transparency of intention: linguistic perspective taking in text. *Cogn. Psychol.* 26:165–208
Keysar B. 1996. Language users as problem solvers. Just what ambiguity problem do they solve? In *Social and Cognitive Psychological Approaches to Interpersonal Communication*, ed. SR Fussell, RJ Kreuz. Mahwah, NJ: Erlbaum. In press
Kintsch W. 1988. The role of knowledge in discourse comprehension: a constructive-integration model. *Psychol. Rev.* 95:163–82
Kintsch W. 1994. Text comprehension, memory, and learning. *Am. Psychol.* 49:294–303
Kintsch W, van Dijk TA. 1978. Toward a model of text comprehension and production. *Psychol. Rev.* 85:363–94
Kintsch W, Welsch D, Schmalhofer F, Zimny S. 1990. Sentence memory: a theoretical analysis. *J. Mem. Lang.* 29:133–59
Klin CM, Myers JL. 1993. Reinstatement of causal information during reading. *J. Exp. Psychol. Learn. Mem. Cogn.* 19:554–60
Kreuz RJ, MacNealy MS. 1996. *Empirical Approaches to Literature and Aesthetics.* Norwood, NJ: Ablex
Levelt WJM. 1989. *Speaking: From Attention to Articulation*. Cambridge, MA: MIT Press
Lewin K. 1951. *Field Theory in Social Science.* New York: Harper & Row
Lightman A. 1993. *Einstein's Dreams.* New York: Warner
Long DL, Golding JM, Graesser AC. 1992. The generation of goal-related inferences during narrative comprehension. *J. Mem. Lang.* 5:634–47
Long DL, Oppy BJ, Seely MR. 1996. Individual differences in sentence- and text-level representations. *J. Mem. Lang.* In press
Lorch RF, Klusewitz MA, Lorch EP. 1995. Distinctions among reading situations. See Lorch & O'Brien 1995, pp. 375–98
Lorch RF, O'Brien JD, eds. 1995. *Sources of Coherence in Reading*. Mahwah, NJ: Erlbaum
MacDonald MC, Pearlmutter NJ, Seidenberg MS. 1994. The lexical nature of syntactic ambiguity resolution. *Psychol. Rev.* 101: 676–703
Magliano JP, Baggett WB, Johnson BK, Graesser AC. 1993. The time course of generating causal antecedent and causal consequence inferences. *Discourse Process.* 16:35–53
Magliano JP, Graesser AC. 1991. A three-pronged method for studying inference generation in literary text. *Poetics* 20:193–232
Mann WC, Thompson SA. 1986. Relational propositions in discourse. *Discourse Process.* 9:57–90
Mannes S. 1994. Strategic processing of text. *J. Educ. Psychol.* 86:577–88
Marslen-Wilson WD, Tyler LK, Koster C. 1993. Integrative processes in utterance resolution. *J. Mem. Lang.* 32:647–66
Mauner G, Tanenhaus MK, Carlson GN. 1995. Implicit arguments in sentence processing. *J. Mem. Lang.* 34:357–82
McClelland JL, Rumelhart DE, eds. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition.* Cambridge, MA: MIT Press
McKoon G, Gerrig RJ, Greene SB. 1996. Pronoun resolution without pronouns: some consequences of memory-based text processing. *J. Exp. Psychol. Learn. Mem. Cogn.* In press
McKoon G, Ratcliff R. 1992. Inference during reading. *Psychol. Rev.* 99:440–66
McNamara DS, Kintsch E, Songer NB, Kintsch W. 1995. Text coherence, background knowledge and levels of understanding in learning from text. *Cogn. Instr.* In press
Miall DS, Kuiken D. 1994. Beyond text theory: understanding literary response. *Discourse Process.* 17:337–52
Millis KK, Graesser AC. 1994. The time-course of constructing knowledge-based inferences for scientific texts. *J. Mem. Lang.* 33:583–99
Millis KK, Graesser AC, Haberlandt K. 1993. The impact of connectives on the memory for expository texts. *Appl. Cogn. Psychol.* 7:317–39
Millis KK, Just MA. 1994. The influence of connectives on sentence comprehension. *J. Mem. Lang.* 33:128–47
Moffett J, McElheny KR. 1995. *Points of View: An Anthology of Short Stories.* New York: Penguin
Morris RK. 1994. Lexical and message-level sentence context effects on fixation times in reading. *J. Exp. Psychol. Learn. Mem. Cogn.* 20:92–103
Morrow DG, Greenspan SL, Bower GH. 1987. Accessibility and situation models in narra-

tive comprehension. *J. Mem. Lang.* 26: 165–87
Murray JD. 1995. Logical connectives and local coherence. See Lorch & O'Brien 1995, pp. 107–25
Myers JL, O'Brien EJ, Albrecht JE, Mason RA. 1994. Maintaining global coherence during reading. *J. Exp. Psychol. Learn. Mem. Cogn.* 20:876–86
Nystrand M. 1986. *The Structure of Written Communication: Studies in Reciprocity Between Readers and Writers.* Norwood, NJ: Ablex
O'Brien EJ, Albrecht JE. 1992. Comprehension strategies in the development of a mental model. *J. Exp. Psychol. Learn. Mem. Cogn.* 18:777–84
O'Brien EJ, Albrecht JE, Hakala CM, Rizzella ML. 1995. Activation and suppression of antecedents during reinstatement. *J. Exp. Psychol. Learn. Mem. Cogn.* 21:626–34
Ohtsuka K, Brewer WF. 1992. Discourse organization in the comprehension of temporal order. *Discourse Process.* 15:317–36
Otero J, Kintsch W. 1992. Failures to detect contradictions in a text: what readers believe versus what they read. *Psychol. Sci.* 3:229–35
Perfetti CA, Britt MA. 1995. Where do propositions come from? In *Discourse Comprehension: Essays in Honor of Walter Kintsch*, ed. CA Weaver, S Mannes, CR Fletcher, pp. 11–34. Hillsdale, NJ: Erlbaum
Perfetti CA, Britt MA, Georgi MA. 1995. *Text-Based Learning and Reasoning: Studies in History*. Hillsdale, NJ: Erlbaum
Perfetti CA, Marron MA, Foltz PW. 1996. Sources of comprehension failure: theoretical perspectives and case studies. In *Reading Comprehension Disabilities*, ed. C Cornoldi, J Oakhill. Hillsdale, NJ: Erlbaum. In press
Potts GR, Keenan JM, Golding JM. 1988. Assessing the occurrence of elaborative inferences: lexical decision versus naming. *J. Mem. Lang.* 27:399–415
Pressley M, Symons S, McDaniel MA, Snyder BL, Turnure JE. 1988. Elaborative interrogation facilitates acquisition of confusing facts. *J. Educ. Psychol.* 80:268–78
Rayner K, Garrod S, Perfetti CA. 1992. Discourse inferences during parsing are delayed. *Cognition* 45:109–39
Rayner K, Pacht JM, Duffy SA. 1994. Effects of prior encounter and global discourse bias on the processing of lexically ambiguous words: evidence from eye fixations. *J. Mem. Lang.* 33:527–44
Rinck M, Williams P, Bower GH, Becker ES. 1996. Spatial situation models and narrative understanding: some generalizations and extensions. *Discourse Process.* 21:23–56
Rubin DC. 1995. *Memory in Oral Traditions: The Cognitive Psychology of Epic, Ballads, and Counting-Out Rhymes.* New York: Oxford Univ. Press
Sanders TJM, Spooren WPM, Noordman LGM. 1992. Toward a taxonomy of coherence relations. *Discourse Process.* 15:1–36
Sanford AJ, Garrod SC. 1981. *Understanding Written Language: Explorations in Comprehension Beyond the Sentence.* New York: Wiley
Schmalhofer F, Glavanov D. 1986. Three components of understanding a programmer's manual: verbatim, propositional, and situational representations. *J. Mem. Lang.* 25: 279–94
Schober MF. 1995. Speakers, addressees, and frames of reference: whose is minimized in conversations about locations? *Discourse Process.* 20:219–47
Seifert CM, McKoon G, Abelson RP, Ratcliff R. 1986. Memory connections between thematically similar episodes. *J. Exp. Psychol. Learn. Mem. Cogn.* 12:220–31
Singer M, Graesser AC, Trabasso T. 1994. Minimal or global inference during reading. *J. Mem. Lang.* 33:421–41
Singer M, Halldorson M, Lear JC, Andrusiak P. 1992. Validation of causal bridging inferences in discourse understanding. *J. Mem. Lang.* 31:507–24
Singer M, Ritchot K. 1996. The role of working memory capacity and knowledge access in text inference processing. *Mem. Cogn.* In press
St. John MF. 1992. The story Gestalt: a model of knowledge-intensive process in text comprehension. *Cogn. Sci.* 16:271–306
Stein NL, Liwag MD. 1996. Children's understanding, evaluation, and memory for emotional events. In *Developmental Spans in Event Comprehension and Representation*, ed. P van den Broek, P Bauer, T Bourg. Hillsdale, NJ: Erlbaum. In press
Suh SY, Trabasso T. 1993. Inferences during reading: converging evidence from discourse analysis, talk-aloud protocols, and recognition priming. *J. Mem. Lang.* 32: 279–300
Trabasso T, Magliano JP. 1996. Conscious understanding during comprehension. *Discourse Process.* In press
Trabasso T, van den Broek P. 1985. Causal thinking and the representation of narrative events. *J. Mem. Lang.* 24:612–30
van den Broek P. 1990. Causal inferences and the comprehension of narrative texts. See Graesser & Bower, pp. 25:175–94
van den Broek P, Lorch RF. 1993. Network representations of causal relations in memory for narrative texts: evidence from primed recognition. *Discourse Process.* 17: 75–98

van Dijk TA, Kintsch W. 1983. *Strategies of Discourse Comprehension.* New York: Academic

van Oostendorp H, Zwaan RA, eds. 1994. *Naturalistic Text Comprehension.* Norwood, NJ: Ablex

Voss JF, Silfies LN. 1996. Learning from history text: the interaction of knowledge and comprehension skill with text structure. *Cogn. Instr.* In press

Whitney P, Budd D, Bramucci RS, Crane RS. 1995. On babies, bathwater, and schemata: a reconsideration of top-down processes in comprehension. *Discourse Process.* 20: 135–66

Whitney P, Ritchie BG, Clark MB. 1991. Working memory capacity and the use of elaborative inferences in text comprehension. *Discourse Process.* 14:133–45

Zwaan RA. 1994. Effects of genre expectations on text comprehension. *J. Exp. Psychol. Learn. Mem. Cogn.* 20:920–33

Zwaan RA. 1996. Processing narrative time shifts. *J. Exp. Psychol. Learn. Mem. Cogn.* In press

Zwaan RA, Brown CM. 1996. The influence of language proficiency and comprehension skill on situation model construction. *Discourse Process.* In press

Zwaan RA, Langston MC, Graesser AC. 1995a. The construction of situation models in narrative comprehension: an event-indexing model. *Psychol. Sci.* 6:292–97

Zwaan RA, Magliano JP, Graesser AC. 1995b. Dimensions of situation model construction in narrative comprehension. *J. Exp. Psychol. Learn. Mem. Cogn.* 21:386–97