SITUATION SCHEMAS, THEMATIC ROLES AND GRAMMATICAL MORPHEMES

By

Todd R. Ferretti

Graduate Program
in
Psychology

Submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy

Faculty of Graduate Studies
The University of Western Ontario
London, Ontario
July, 2000

Abstract and Keywords

This research investigated how people combine morpho-syntactic information (e.g., aspect) with world knowledge of situations when they read verbs and noun phrases in isolation. In Experiment 1, subjects read verb phrases presented for a brief duration (250 ms), that were marked with either imperfect (*was verbing*) or perfect aspect (*had verbed*). They then named visually presented targets that were either typical locations (*was skating - arena*) or instruments (*was stabbing - dagger*). Typical locations of events were more highly activated when the verbs referenced the situations as ongoing (imperfective) versus completed (perfect). The instrument results were less clear, however. These results indicate that verb aspect can modulate the activation of world knowledge about typical locations of situations quickly when verbs are read in isolation.

These results were extended in Experiments 2, 3, and 4 by examining how people integrate world knowledge of agents and patients in situations with the aspectual properties of present and past participles to constrain interpretation of isolated phrases such as *arresting cop* and *arrested crook*. An implemented competition model was used to generate predictions about how people interpret these types of phrases. The model correctly predicted that subjects combined typical patients more easily with past participles (*arrested crook*) than with present participles (*arresting crook*). Interestingly, they often interpreted phrases like *arresting crook* as verb phrases when the head noun was a great patient / terrible agent. Furthermore, subjects combined typical agents with present participles (*arresting cop*) more easily than with past participles (*arrested cop*). Thus the activation of world knowledge of event participants is modulated by grammatical morphemes, and people equally weight these sources of

information when combining them to constrain thematic role assignment during on-line interpretation of phrases.

Key words: aspect, thematic roles, situation schemas, noun phrase interpretation, competition modeling

Acknowledgements

First and foremost, I would like to thank my advisor, Dr. Ken McRae, for his excellent direction and advice, and for the time he invested in making editorial corrections.

Finally, I would like to thank Anneke Olthof and my parents for their love and support.

TABLE OF CONTENT Page

LIST OF TABLES

LIST OF FIGURES

LIST OF APPENDICES

Chapter 1: Introduction

Language comprehension involves the rapid integration of various types of semantic and syntactic knowledge (Altmann, Garnham, & Dennis, 1992; Spivey-Knowlton & Sedivy, 1995; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995). The fact that people comprehend language quickly and incrementally suggests that our knowledge structures are organized so that many types of information can be computed and used quickly in an integrated manner (Marslen-Wilson, 1975). One component of language comprehension for which the integrative nature of processing has been highlighted is that of assigning nouns to a verb's thematic roles. One way to facilitate this integration is if a verb makes available information about its syntactic arguments and, in addition, the semantic features specific to the entities and objects that typically fill its various thematic roles. The latter process might be viewed in terms of a verb providing immediate access to the generalized situation structure (which could be realized as a schema) that corresponds to the situations/events to which it refers.

The goal of the present thesis is to investigate whether reading verbs and participles marked with grammatical morphemes modulates the activation of thematic role information regarding the typical locations, instruments, agents, and patients that are part of the situations to which they refer. In line with this goal, a perspective is developed that draws on the literatures concerning thematic role processing, grammatical morphemes, conceptual combination, sentence comprehension, schema theory, semantic memory, semantic priming, autobiographic memories, and inferencing.

Chapter 1 provides a review of the relevant concepts from some of these areas of research, specifically, those related to schemas and McRae and colleagues' (Ferretti, McRae, & Hatherell, in press; McRae, Ferretti, & Amyote, 1997) theoretical claims regarding the conceptual content of thematic roles. This background is the primary source of motivation for performing the research reported in Chapters 2 and 3. Other research is discussed as it becomes relevant.

1.1 Schemas

Understanding language involves constructing a mental representation that reflects the comprehender's interpretation of the individual situations described in the discourse, and the relationships among them (Johnson-Laird, 1983; Kintsch, 1988; Rumelhart, 1980; 1991; van Dijk & Kintsch, 1983; Zwaan, Langston, Graesser, 1995). There is a general consensus that comprehenders interpret situations and their relations by applying knowledge structures that capture common properties of their past experiences (Goldberg, 1998; MacWhinney, 1999). Researchers have proposed a number of mental structures designed to capture people's world knowledge about situations, with the most common forms being schemas (Rumelhart, 1980; Rumelhart & Ortony, 1977), scripts (or event-memory organization packets, Kolodner, 1983; Shank & Abelson, 1977), and frames (Minsky, 1975).

According to the schema view, knowledge is packaged into units, and embedded in these units is the knowledge people have about the world, including information about objects, situations, and sequences of situations. In most versions of schema theory, situations are organized into structured lists of slots (variables), each of which can contain fillers

(values). The slots represent the general dimensions of a concept along with their default values. For example, the schema corresponding to arresting situations would include a slot that specifies the agent or "arrester", which may have a range of values such as "police", "soldier", and "security guard". In some versions of schema theory, the default value for the "arrester" slot would be "police" because it is the most frequent. According to early perspectives, schemas and their corresponding defaults were viewed as highly-structured information stored in memory and accessed directly from it. Many researchers have since recognized the inadequacy of these views, and have suggested more flexible mechanisms to specify how schemas are organized and used on-line (Rumelhart, Smolensky, McClelland, & Hinton, 1986). Such proposals include connectionist approaches to knowledge representation (McClelland & Rumelhart, 1985) that are related to some instance-based or exemplar views (Hintzman, 1986). The major insight of this newer perspective is that schemas are not explicit entities stored in memory and directly accessed from it, rather, they emerge or are created on-line in the context of the situation or utterance currently being interpreted.

The organization of situations in memory has played an important part in constraining theories of autobiographical memories (Anderson & Conway, 1997; Barsalou, 1988; Brown & Schopflocher, 1998; Lancaster & Barsalou, 1997; Reiser, Black, & Abelson, 1985). Of particular interest has been the level of abstraction that is most likely to capture information encoded in memory about our experiences. For example, a situation such as "eating at a restaurant" can be viewed as a superordinate concept that includes subordinate levels capturing generalized situations, such as "eating", "paying" and "ordering". In addition, generalized situations may

have a subordinate level of representation, such as "lifting a fork" (Reiser et al., 1985). The mid-level of abstraction captured by generalized situations can be viewed as corresponding to the basic level that has been shown by numerous researchers to have a privileged status in cognition (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976; Tversky & Hemenway, 1983). Regardless of controversies surrounding the basic level, researchers agree that generalized situation schemas are important for organizing people's experiences, and that these structures can include detailed and specific knowledge.

The central function of schemas in language comprehension is well-documented in the literature examining how people understand narratives. This research has demonstrated that one of the major functions of a schema is to provide specific information that allows a comprehender to generate predictions about situations referred to in a discourse. For example, researchers have shown that our knowledge about the common properties of situations influences the errors or intrusions people make when recalling a story (Bower, Black, & Turner, 1979), and which information is most likely to be recalled and recognized (Graesser, Woll, Kowalski, & Smith, 1980; Pichert & Anderson, 1977). More recently, it has been demonstrated that this information is a basis for drawing causal relations among situations mentioned in a discourse (Fletcher & Bloom, 1988; Golden & Rumelhart, 1993). Furthermore, under limited conditions, people use this knowledge to elaborate situations during on-line discourse comprehension, such as making elaborative inferences about common instruments (Dosher & Corbett, 1982; Lucas, Tanenhaus, & Carlson, 1990; McKoon & Ratcliff, 1981), and making predictions about probable outcomes (Keefe & McDaniel, 1993; McKoon & Ratcliff, 1986; Potts, Keenan, &

Golding, 1988). Finally, in the sentence processing literature, detailed, specific world knowledge about situations has been shown to be a key component of syntactic ambiguity resolution (McRae, Spivey-Knowlton, & Tanenhaus, 1998; Schmauder & Egan, 1998), expectancy generation in sentence processing (Taraban & McClelland, 1988), and lexical ambiguity resolution (Till, Mross, & Kintsch, 1988; Vu, Kellas, & Paul, 1998).

The foregoing discussion illustrates the important function that generalized situation knowledge plays during language comprehension. However, despite numerous demonstrations that it is an essential aspect of comprehension, scant research has integrated it with a proposal regarding the conceptual information made available by verbs about situations and their common properties, and not surprisingly, there has been even less research on how this information may be constrained by grammatical morphemes that reference situations from various temporal viewpoints. Clearly, determining the level of detail of this information and how it interacts with grammatical morphemes is an important step toward understanding the on-line availability of knowledge about situations. One of the main reasons for this lack of synthesis is that thematic role knowledge has traditionally been viewed as distinct from world knowledge about the situations they denote, particularly in any theory that has emerged from the Chomskyan tradition (e.g., Caplan, Hildebrandt, & Waters, 1994; Schlesinger, 1995; but see Johnson-Laird, 1983, and McCawley, 1968, for a contrasting view). Thus, research examining thematic role information has tended to focus on linguistically or syntactically relevant information, whereas research on schemas has tended to ignore thematic role knowledge and assignment. The following section discusses a recent view that suggests that the information that must be deployed to use a verb

correctly includes specific world knowledge about the common characteristics of situations that they denote.

1.2 Verb-Specific Concepts

Recently, McRae et al. (1997) incorporated and extended the important work of Carlson and Tanenhaus (1988), Dowty (1991), Fillmore (1968), Gentner (1981), Jackendoff (1983), Langacker (1987), and others to construct a theory of thematic roles that incorporates situation-specific information. McRae et al. focused on the fact that an important aspect of verb meaning and situation structure is the relations among the entities that commonly participate, a type of information that might be instantiated in schemas. In their view, each thematic role of each verb is a concept that is formed through everyday experiences during which people learn about who and what play specific roles in specific situations. This knowledge then serves as an important constraint in on-line language comprehension and production. For example, a person's representation of the agent role of *accuse* results from her experiences with people who accuse others in everyday episodes and in linguistic descriptions of them. They further speculated that this knowledge is computed immediately as a consequence of hearing or reading the verb *accuse*, and thus is a key factor in driving on-line thematic assignment.

McRae et al. (1997) illustrated that people possess this knowledge and can verbalize at least a portion of it by showing that subjects easily produce features for agent and patient concepts such as "someone who accuses others" or "someone who is accused". In contrast to general features such as <cause>, <volitional involvement>, and <sentient> on which Dowty (1991) and Schlesinger (1995) have focused, the subject-generated features were

quite specific in nature. For example, for "someone who accuses others", subjects provided features such as <is mean>, <is judgmental>, and <is knowledgeable>. Note that as in prototype theory, and contrary to a basic premise of selectional restrictions, these features are characteristic, rather than defining. McRae et al. used the featural descriptions to demonstrate that thematic roles have internal structure, analogous to artifact and living thing noun concepts (Rosch & Mervis, 1975). To do this, a second group of subjects rated the importance of the features to specific noun concepts, such as *witness* (e.g., the importance of <is knowledgeable> to the concept *witness*). An estimate of role/filler featural similarity was calculated as the mean featural importance rating across that thematic role's features. A third group of subjects rated the typicality of nouns with respect to specific thematic roles, as in, "On a scale of 1 to 7, how common is it for a witness to accuse someone?" Role/filler featural similarity was a strong predictor of role/filler typicality ratings, suggesting that thematic relations have internal structure and thus are represented in a similar manner as nominal concepts. Furthermore, their research indicated that the internal structure is captured by specific world knowledge such as <is authoritative>.

These experiments do not, however, speak to the issue of whether this knowledge is tied tightly to a verb or is better considered as part of a distinct episodic or semantic memory system. McRae et al. (1997) used theoretical notions from the models of McClelland and Rumelhart (1985) and Hintzman (1986) to suggest that, in terms of on-line processing, it is tied tightly to a verb although it derives from episodic traces. (Note that this implies a somewhat broader view of lexical semantic memory than is often taken in that there are no clear distinctions among episodic memory,

autobiographical memories, and semantic memory, distinctions that were popularized by Tulving, 1983. See McKoon, Ratcliff, & Dell, 1986, for particularly forceful arguments against Tulving's distinction.). In these models, specific episodes are encoded as traces laid down in an independent or superimposed form. When a word is read or heard, the resultant computed representation is akin to a feature-based prototype that is modulated by the quantity and quality of the context included in the memory probe. Thus, a computed thematic role concept can be viewed as a set of semantic features that constitute the most frequent and/or intercorrelated features of the noun concepts that typically fill that role (shaded to various degrees by context). For example, most people have witnessed or participated in multiple accusing situations. Individuals who fill the agent role of *accuse* often possess certain features, and these are most likely computed as part of the agent role representation of the verb. That is, the hypothesis is not that a list of possible nouns is computed; rather, a distributed featural-type of representation is computed that overlaps more or less with various lexical concepts. McRae et al.'s (1997) view of the organization of the conceptual aspects of thematic role knowledge draws on connectionist views of schemas (Golden & Rumelhart, 1993; Rumelhart et al., 1986; Sharkey & Sharkey, 1992). This perspective entails that conceptual thematic role knowledge emerges from the correlations among the knowledge structures represented in memory. Furthermore, this kind of representational system can utilize immediately both specific and general information that people possess about the world, and is highly sensitive to contextual influences.

Ferretti et al. (in press) provided further evidence for this view by demonstrating that verbs with well-defined roles make available event-

specific information about their common role fillers. Specifically, in a series of short SOA priming experiments they illustrated that verbs made available information immediately about their common agents, (*arresting-cop*), patients (*arresting-crook)*, patient features (*manipulating-naive*) and instruments (*stirred-spoon*). However, they did not find that verbs made available information about their typical locations (*swam-ocean*).

They also directly examined whether it is reasonable to consider the primed information as thematic role knowledge by testing whether syntactic cues modulate activation of common participants. Verbs were presented as part of short auditorally-presented active (*She arrested the*) and passive sentence fragments (*She was arrested by the*), and subjects named a visual target (*cop*) presented at their offset. Agents and patients were presented in their congruent role (*She was arrested by the / cop*) and their incongruent role (*She arrested the / cop*). Priming was found only when nouns were in their congruent roles, indicating that situation-specific (verb-specific) information of this sort should be considered part of thematic role knowledge. This experiment is also important to the current discussion because it demonstrates how the grammatical category of *voice* (i.e., active and passive forms of verbs) can modulate activation of world knowledge of events; when verbs were presented in their passive form, activation of knowledge about the verb's common role fillers was constrained to only the participant filling the expected role of the verb, the agent, and further that this process was completed rapidly following the presentation of the verb. These results suggest that grammatical morphemes may play a significant role in modulating the activation of world knowledge about situations.

1.3 Overview

The purpose of the present thesis is to examine how aspectual morphemes constrain activation of world knowledge of situations. These grammatical morphemes temporally reference the development of situations from various viewpoints (e.g., as ongoing or completed). The question addressed in this thesis is whether the viewpoints afforded by these grammatical morphemes can lead to the foregrounding and backgrounding of world knowledge concerning entities, objects, and locations of situations during on-line language comprehension.

To understand why these morphemes may play a crucial role in modulating world knowledge, one has to consider both the causal and temporal structure of events because they are closely intertwined (see Figure 1). For example, the temporal structure of events can be decomposed into 3 general components, including a beginning, middle, and end. These temporal components correspond closely to the causal properties of events, which include initiating conditions (beginning), the actual event (middle), and the resultant states (end) (Moens & Steedman, 1988). Related to these causal components are the entities, objects, and locations that typically are part of events. Importantly, the entities and objects in events will be more or less salient in the different causal components of events. As illustrated in Figure 2, agents typically are associated with the initiating conditions of events because they tend to cause them to occur, and patients are typically associated with the resultant states of the event because they typically undergo a change of state as a result of the action performed by some agent. Instruments tend to be part of the initiating conditions because they are used by some agent to perform some action. Location and instrument

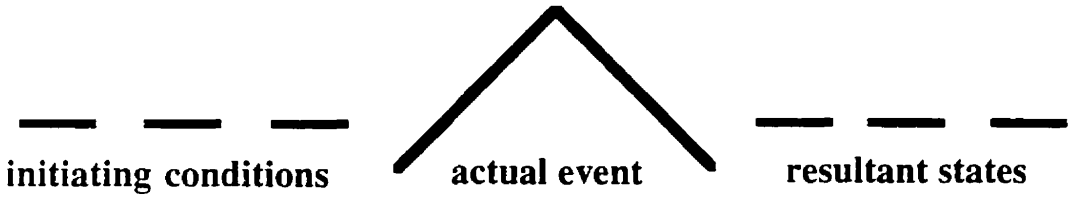

Figure 1. *The temporal and causal structure of events.*

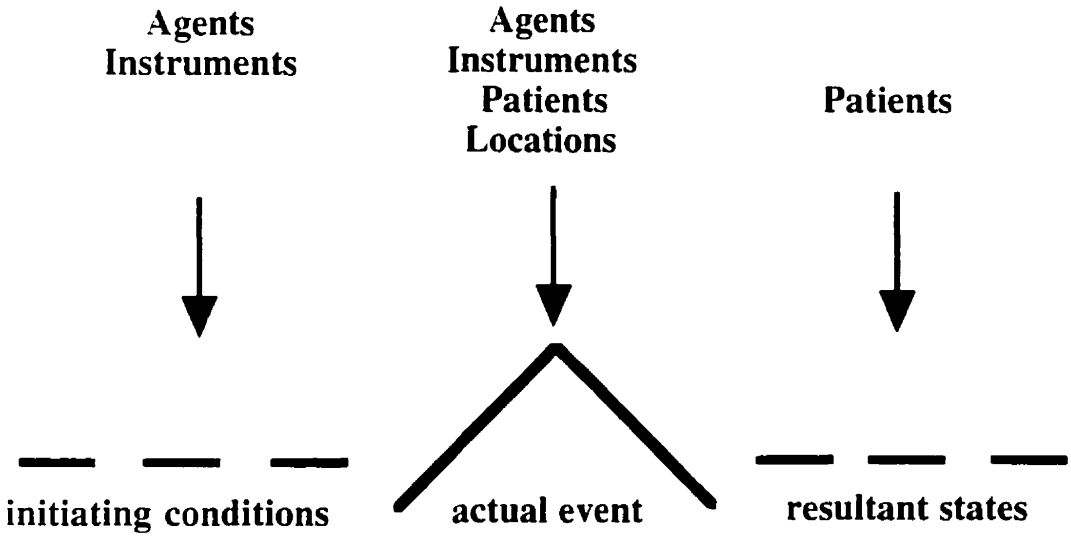

Figure 2. *Thematic roles and their general relationship with the temporal and causal structure of events.*

information is present when the situations are ongoing, but not when they are completed.

Chapter 2 examines the role of verb-aspect in modulating activation of world knowledge about the typical locations and instruments of specific situations using a priming methodology. Chapter 3 extends these results in a number of important ways. First, it investigates how people actively combine morpho-syntactic information with world knowledge to constrain thematic role assignment. Second it extends the results on processing of verb-phrases to noun phrase interpretation. Third, a norming study is presented that examines how people interpret these NPs off-line. This information is used subsequently to motivate a competition model that simulates the combinatorial process involved in the on-line experiments (Experiment 3 and 4). The modeling makes the proposed combinatorial process explicit, and leads to testable predictions regarding how people interpret noun phrases. Finally, in Chapter 4, the implications of the results presented in Chapters 2 and 3 are discussed. This thesis makes a significant contribution to the literature by demonstrating how the temporal and causal structure of situations interacts with grammatical morphemes to constrain activation of world knowledge of entities, objects, and locations in events during on-line language comprehension.

Chapter 2: Aspect and Situation Structure

The grammatical category of verb-aspect has been shown to be important for foregrounding and backgrounding of entities and situations mentioned in a text (e.g., Carreiras, Carriedo, Alonso, & Fernandez, 1997; Morrow, 1985a; 1990; Morrow, Greenspan, & Bower, 1987). Despite this, scant research has examined how information not explicitly mentioned in text, but used to drive understanding (i.e., world knowledge about events), is foregrounded and backgrounded by different aspect morphemes. Situations unfold over time and depending on how aspectual morphemes reference the temporal development of the situation (e.g., beginning, middle, end), one might expect to find conceptual properties of situations consistent with the temporal reference to be more salient.

Chapter 2 begins with a review of the properties of the dynamical situations and aspectual categories used in the following research. The main purpose of this section is to demonstrate how imperfective and perfect aspect are sensitive to the temporal and causal contours of dynamical situations. Understanding the interaction between aspect and these situations is important for making predictions about how these grammatical morphemes modulate activation of world knowledge about typical locations and instruments in Experiment 1, and for demonstrating how the present and past forms of participles constrain activation of information about typical agents and patients in Experiments 2 to 5. In the second part of Chapter 2, the focus shifts to recent psycholinguistic studies examining the role of aspect in constraining discourse model construction. In the final part of the chapter, an on-line experiment is presented that provides insight into how aspect modulates the early activation of world knowledge about common locations and instruments in situations.

2.1 Dynamic Situations and Aspect

The situations that are the focus of the following research are typically classified as dynamical situations because they involve change over time. When a situation involves a single change of state they are classified as *events* (e.g., winning), and those that involve multiple changes of state are considered *processes*. Processes can be progressive (e.g., lift) or cyclical (walk) (Lyons, 1977). In general, dynamical situations are viewed as involving a period that leads up to the actual change of state, a period in which the change(s) of states occur, and a period that follows. In terms of the causal structure of situations, the period that preceeds the changes of state capture the preparatory or initiating processes by which the situation is accomplished, and the period following the situation captures the consequent or resultant states (Moens & Steedman, 1988; Nakhimovsky; 1988). Figure 1 illustrates these components of a situation.

The causal and temporal structure of dynamical situations make them useful for the present purpose because they involve subcomponents that can be specified more directly by aspectual categories. The two aspectual categories examined below are *imperfective* and *perfect* aspect.

Imperfective aspect makes specific reference to the internal structure of situations by focusing on their ongoing development, but does not make reference to their terminative phases (see Figure 3). The imperfective category is frequently subdivided by linguists into progressive and habitual categories. The English progressive is marked with the verbal form *be - ing* (e.g., is scoring, was scoring, will be scoring). The progressive form focuses on the internal phases of dynamical situations, and the habitual form differs from the progressive form in that it refers to a *series* of a particular event, or the *iteration* of a process several times (Comrie, 1976; Smith, 1991).

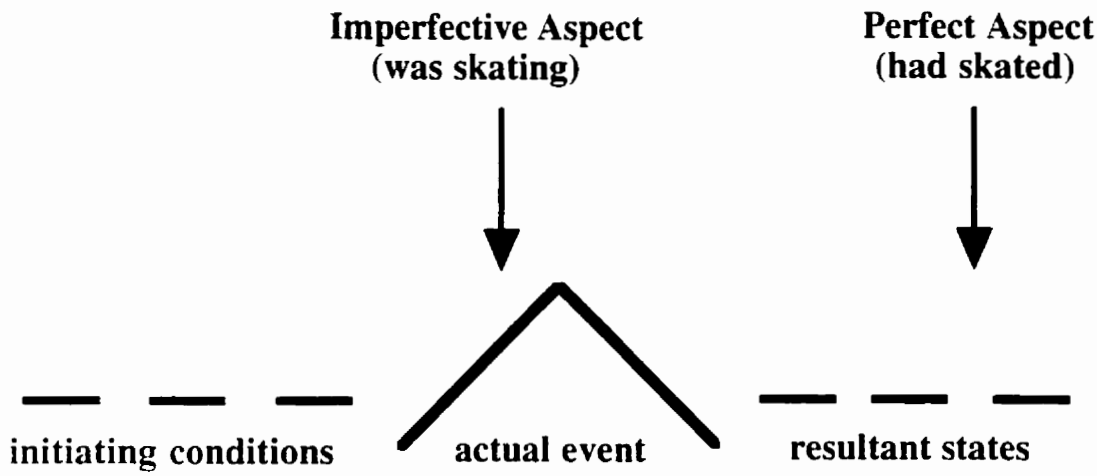

Figure 3. *Referencing the temporal and causal structure of events with imperfective and perfect aspect.*

The perfect aspect refers to times later than the situation and places emphasis on the resultant phases (or states) rather than on inceptive phases (Comrie, 1976; see Figure 3). In English this aspectual category is marked with *to have* + *past participle* (e.g., has scored, had scored, will have scored). In general, the perfect functions to indicate the continuing relevance of a past situation into the present or to other reference times (Comrie, 1976).

2.2 Psycholinguistics and Aspect

Research in psycholinguistics has only begun to examine the role that aspect plays in language comprehension. The few studies that have examined it have been concerned primarily with how people use this information to build a mental model of a text. These results add to a growing list of factors that have been demonstrated to constrain mental model construction (Johnson-Laird, 1983; Kintsh, 1988; Morrow, 1985a; 1990; Morrow, Greenspan, & Bower, 1987; Sanford & Garrod, 1981; van Dijk & Kintsch, 1983; Zwaan, Langston, & Graesser, 1995).

Evidence for the psychological relevance of these distinctions has been explored in a recent series of studies by Morrow and colleagues (Morrow, 1985a; 1990). The main goal of these studies was to examine how grammatical morphemes such as prepositions and verb-aspect constrain construction of spatially organized mental models. For instance, Morrow (1990) examined how readers combine these morphemes to infer the location of a figure moving through a house. Subjects memorized a layout of a house that consisted of a floor plan that illustrated the location of a kitchen, bathroom, living room, and hallway. This methodology allowed for subjects to potentially be specific about the location of the figure in the

situations described by these sentences. Subjects read sentences such as examples (2.1) and (2.2), and then were instructed to indicate where they thought the figure (e.g., John) was located at the time described by the sentence.

(2.1) John was walking from the kitchen through the living room to the bedroom.

(2.2) John walked from the kitchen through the living room to the bedroom.

Depending on the nature of the prepositions and the aspectual morphemes used in the sentences, Morrow was able to manipulate the reader's spatial model of the specific location of the figure moving through a house. Morrow found that when the imperfect aspect was used (e.g., 2.1) readers were more likely to locate the figure somewhere on an unmentioned path. Alternatively, when the verbs were in their perfective form, readers consistently located the figure at or inside the goal (2.2).[1] Thus readers used the combination of aspect and the information provided by prepositions to infer the location of the figure in a spatially organized mental model. Furthermore, these results shed light on how comprehenders use aspect to make salient various components of events.

The ability of aspectual morphemes to influence the construction of a mental model via the foregrounding of information about multiple characters and events has also been examined (Carreiras et al., 1997; Morrow, 1985b; Morrow et al., 1987). Carreiras et al. (1997) provide a good example of the role that verbs presented in their imperfective and perfect forms can play in this process. Subjects read narratives such as (2.3).

---

[1]Perfective aspect is differentiated from the other two categories by referencing the whole event and the different phases or parts that make up an event are ignored (Smith, 1991).

(2.3) John works as a waiter in a restaurant.

Mary eats there every day.

a. John was finishing his shift...

b. John had finished his shift...

when Mary arrived at the restaurant.

c. She asked for the dish of the day.

d. She read the newspaper while waiting for the food.

Subjects read either the version with John and Mary simultaneously involved in ongoing situations (2.3a) or at different times (2.3b). The rational was that John should be more in focus, and thus more accessible, when he is presented as participating in an action that was ongoing when Mary arrived. Under these circumstances, both characters are in focus. Alternatively, describing John with a verb marked with perfect aspect should function to background John and foreground Mary in the discourse model (i.e., make John less accessible).

Carreiras et al. (1997) examined the ability of readers to use this temporal information to constrain their mental model of the text by probing for activation of the character included in the third sentence (i.e., John). The character was presented either immediately following the sentence mentioning the two actions (2.3b), after one filler sentence (2.3c), or after two filler sentences (2.3d). Subjects were instructed to respond as quickly as they could as to whether they thought the word had appeared in the previous text (i.e., a recognition task). Subjects were slower to recognize the target when it was presented as occurring at a different point in time in regard to the action of the other character, than when it was presented as occurring simultaneously with the character. Moreover, the influence of this manipulation was evident immediately following the sentence

describing the two actions, and the advantage in accessibility remained after the first filler sentence. This study provides additional support for the importance of aspect for constraining the development of mental models during on-line processing. The temporal relations between characters presented in a narrative clearly play an important role in constraining a reader's developing discourse model by modulating which characters and situations are in focus. The strength of this variable is evidenced by the length of the sustained activation for the characters in focus.

2.3 Modulating Situation Knowledge with Aspect

One question that has received little attention is how well-known knowledge of situations is constrained by aspect. Only a recent study by Truit and Zwaan (1997) provides some evidence of this possibility. They demonstrated that readers inferred implied instruments of situations, and maintained activation of these instruments when the verbs were in their imperfective (e.g., He was pounding the nail.) but not their perfective form (He pounded the nail.). Specifically, readers took longer to recognize that a implied instrument was not in the previous paragraph, and named the implied instrument more quickly, when the critical situation was viewed as ongoing versus completed. Thus, accessibility of implied instruments is facilitated by the imperfective form of a verb. Unfortunately, this study did not probe for activation of instruments immediately following the verb. Thus, it is not clear when the information about instruments became *available* during on-line processing.

Experiment 1 furthers this line of research in a few notable ways. First, it probed for activation of thematic role information immediately after verb-phrases so that early activation of thematic role information was

indexed for different aspect morphemes. Second, the experiment examined whether aspect modulates the activation of the typical locations of which particular events occur. Finally, the experiment contrasted imperfective and perfect aspect. Perfect aspect was used because, unlike perfective aspect that references an entire event as completed, the perfect references only the resultant states. Therefore, the perfect aspect provides a stronger temporal contrast with the ongoing viewpoint afforded by imperfective aspect.

2.4 Experiment 1

The location and instrument pairs chosen for this study were a subset of the pairs used in Ferretti et al. (in press). Recall that they found priming for instrument (*dug-spade*) but not location pairs (*skated-arena*). One explanation for these results is that locations tend to be background information for most situations and thus may not be as salient in the causal structure of events in comparison to instruments, which are typically associated with agents and thus the initiating conditions of events. However, the research reviewed above indicates that backgrounded information about situations can be foregrounded by imperfective aspect (Carreiras et al., 1997). Furthermore, this grammatical morpheme has been shown by Morrow (1990, 1985) to highlight the path along which figures are moving. Taken together, these results suggest that Ferretti et al. may not have obtained priming for the locations because the verb primes were not presented in their imperfective form. That is, the common locations for the situations denoted by their verbs might be activated to a greater degree if the same verbs are marked with imperfective aspect.

Instruments provide an interesting contrast to locations because they are associated more saliently with all aspects of the temporal and causal structure of situations. Specifically, instruments are used by agents to cause some change of state in some other entity or object. Therefore, instruments are associated with the initiating conditions and middle portions of situations because of their close relationship with agents. Thus, if a situation is viewed as ongoing, the instrument is clearly present. However, it is not entirely clear what should happen when verbs are marked with perfect aspect. Recall that one function of perfect aspect is to indicate the continuing relevance of some past event. For events involving instruments (e.g., someone stabbing someone else), the resultant state is heavily constrained by the type of instrument employed. Perhaps this contributes to the continuing relevance of instruments, and thus activation of knowledge about the common instruments used in situations. Another possibility, however, is that because the perfect aspect clearly indexes situations as completed, the instruments in these situations are less likely to be physically present, and thus less likely to be activated.

Method

Subjects

Forty-eight subjects participated for course credit, 12 per list. In all experiments reported in this article, subjects were native English-speaking psychology undergraduates from the University of Western Ontario who had normal or corrected-to-normal visual acuity. Each subject participated in only one of the experiments reported herein.

Materials

Twenty-four verb-location and twenty-four verb-instrument pairs were chosen on the basis of a norming study in which 40 subjects rated how

typical a specific location or instrument is for the situation denoted by a verb (see Appendix A). Subjects provided ratings on a scale of 1 (very uncommon) to 7 (very common) for questions similar to example (2.4) for locations, and (2.5) for instruments.

(2.4) How common is it for someone to skate in each of the following locations?

a backyard _____

a park _____

a pond _____

a rink _____

an arena _____

a lake _____

a schoolyard _____

(2.5) How common is it for someone to use each of the following to perform the action of stirring?

a spoon _____

scissors _____

a fork _____

a straw _____

a stick _____

For the items chosen, the mean typicality rating was high for both the 24 locations (6.7, *range* = 6.0 - 7.0) and the 24 instruments (6.3, *range* = 5.6 - 6.9).

There were 4 lists, each containing 24 verb-location and 24 verb-instrument items. Half of the verbs for both the location and instrument

pairs were marked with imperfective aspect (e.g., *was skating-arena* and *was digging-spade*), and the other half were marked with perfect aspect (*had studied-library* and *had inflated-pump*). In each list, the verb-location and verb-instrument trials were divided into 12 related (e.g., *had skated-arena* and *was digging-spade*) and 12 unrelated trials (e.g., *had prayed-arena* and *had inflated-spade*). The unrelated trials were formed by re-pairing the related verb-noun combinations, with four exceptions. Four of the verbs, *fish*, *hunt*, *gamble*, and *ate*, were used with both a location and an instrument. To ensure that no subject was presented with any word more than once, these verbs were changed to the semantically similar verbs *trawl*, *track*, *bet*, and *devour* for the unrelated trials. Thus each target noun served as its own control, and no subject saw any word more than once. The four lists included the same set of 48 filler trials composed of unrelated verb-noun pairs. Half of the filler trials were composed of verbs marked with imperfective aspect and the other half were marked with perfect aspect. The relatedness proportion was .25 (24/96). Finally, 12 practice trials were constructed that included 6 verbs marked with imperfective aspect and 6 verbs marked with perfect aspect. Moreover, the 12 trials were composed of 2 related verb-location pairs, 2 related verb-instrument pairs, and 8 were unrelated verb-noun pairs.

Procedure

Stimuli were displayed on a 14-inch Sony Trinitron monitor controlled by a Macintosh LC630 using PsyScope (Cohen, MacWhinney, Flatt, & Provost, 1993). Each trial consisted of the following events: a focal point (+) in the center of the screen for 250 ms; the prime (*was digging*) for 250 ms; and then the target (*spade*) until the subject responded. The intertrial interval was 1500 ms. Subjects were instructed to read the verb phrase and

to name the noun into the microphone as quickly and accurately as possible. Latencies were recorded with millisecond accuracy via a microphone attached to a CMU button box that measured the time between the onset of presentation of the target word and the naming response. Testing sessions began with the practice trials. It took approximately 20 minutes to complete the task.

Design

Four-way analyses of variance were conducted using naming latency as the dependent variable. The three independent variables of interest were thematic role (locations vs. instruments), relatedness (related vs. unrelated), and aspect marking (imperfective vs. perfect). Aspect marking and relatedness were within subjects ($F1$) and within items ($F2$). Role was within subjects and between items. Planned comparisons were conducted to investigate priming effects for locations and instruments with both forms of verb-aspect. Finally, list and item rotation group was used as a between factor to stabilize any variance caused by rotating items and subjects across the different lists (Pollatsek & Well, 1995).

Results

Mean naming latencies by condition are presented in Table 1. Note that a trial was excluded from the latency analyses if the response was incorrect. Because pronunciation errors occurred on less than 2% of the trials, they were not further analyzed. In addition, less than 1% of the trials were discarded because of extraneous noise or failure to trigger the microphone. Finally, response latencies greater than three standard deviations above the grand mean were replaced by that value (less than 1% of the scores).

	was verbing		had verbed	
Role Type	*M*	*SE*	*M*	*SE*
Location role				
Unrelated	601	15	597	12
Related	578	12	597	14
Facilitation	23**		0	
Instrument role				
Unrelated	589	12	585	12
Related	577	12	577	13
Facilitation	12*		8	

**Significant by subjects and items
*Significant by subjects

Table 1

Mean naming latencies (ms) for location and instrument roles in Experiment 1.

The three-way interaction among grammatical marking, relatedness, and role was marginally significant ($F1(1, 44) = 3.74$, $p < .06$, $F2(1, 44) = 2.11$, $p > .1$.[2] This interaction occurred because grammatical marking interacted with relatedness for location role, $F1(1, 44) = 6.51$, $F2(1, 22) = 5.72$, but not for the instrument role, $F < 1$ in both analyses. Planned comparisons for the simple interactions revealed that there was a 23 ms priming effect for locations when verbs were marked with imperfective aspect, $F1(1, 44) = 21.31$, $F2(1, 44) = 9.49$, whereas there was no priming effect for locations when the verbs were marked with perfect aspect, $F < 1$ in both analyses. For the instrument role, there was a 12 ms priming effect when verbs were marked with imperfective aspect that was significant by subjects, $F1(1, 44) = 5.62$, but not by items $F2(1, 44) = 2.37$, $p > .1$. There was a nonsignificant 8 ms priming effect for instruments when verbs were marked with perfect aspect, $F1(1, 44) = 2.76$, $p > .1$, $F2(1, 44) = 1.33$, $p > .2$.

The two-way interaction between relatedness and marking was significant in the subject analysis only, $F1(1, 44) = 4.43$, $F2(1, 44) = 2.84$, $p < .1$. The marginal interaction occurred because naming latencies for related items were shorter for imperfective than perfect aspect, $F1(1, 44) = 4.57$, $F2(1, 44) = 3.34$, $p < .08$, whereas the unrelated items were responded to similarly across both types of grammatical markers, $F < 1$ in both analyses. The interaction between role and marking was significant in the subject analysis only, $F1(1, 44) = 5.04$, $F2(1, 44) = 1.05$, $p > .3$. Finally, there was a marginal interaction between relatedness and role, $F1(1, 44) = 4.43$, $F2(1, 44) = 2.84$, $p < .1$.

Collapsed across role and marking, related trials ($M = 582$ ms, $SE = 6.4$ ms) were responded to more quickly than unrelated trials ($M = 593$ ms, $SE =$

[2]For all inferential statistics, $p < .05$ unless otherwise noted.

6.5 ms), $F_1(1, 44) = 17.79$, $F_2(1, 44) = 10.42$. Overall, instrument items (M = 582 ms, SE = 6.3 ms) were responded to more quickly than location items, (M = 593 ms, SE = 6.7 ms), but this was only significant in the subject analyses, $F_1(1, 44) = 16.24$, $F_2 < 1$. The main of effect of marking was not significant, $F_1 < 1$, F_2 (1,44) = 1.22, p > .2

Discussion

The results for location and instrument roles are discussed in turn, followed by a discussion of their implications for theories of word-word priming. A more detailed discussion of the implications of the current results for language comprehension in general is deferred to the General Discussion in Chapter 4 when all the evidence for focusing properties of grammatical morphemes has been presented.

Locations. The results for the location role support previous research showing that the imperfective aspect highlights the path or location of entities in events (e.g., Morrow, 1990). More importantly, they extend previous research by showing that aspect modulates activation of world knowledge of common locations in which events occur. Specifically, subjects named locations more quickly when they were preceded by verbs denoting an event in which they were typical locations than when they were preceded by unrelated verbs, but only when the verbs were marked with imperfective aspect. When the verbs were presented in their perfect form, locations were named no more quickly after related versus unrelated verbs. From an event structure perspective, the probability of an individual staying at the same location after an event is fairly low. For example, if someone had finished swimming, they could be in a house, driving home, or eating in a restaurant. Furthermore, location information is in general less likely to have continuing relevance in the resultant states of events because they

are normally not strongly associated with the causal structure. Finally, these results indicate that Ferretti et al. (in press) did not find activation for typical locations because their verbs referenced events as completed. Similar to Ferretti et al., when the events were viewed as completed there was again no evidence for activation of common locations. However, the same verbs led to the activation of typical locations when the events were referenced as ongoing.

Instruments. Similar to the results for locations, evidence for early activation of knowledge about common instruments (*dagger*) was found when the verb primes were marked with imperfective aspect (*was stabbing*). Although there was some evidence of priming for instruments when they were preceded by the same verbs marked with perfect aspect, the size of the effect was not significant. These results are consistent with previous research by Truit and Zwaan (1997) showing that sentences containing verbs marked with imperfective aspect (The man was pounding the nail) led to the activation of world knowledge about common instruments (hammer) that was not activated when the same events were viewed as completed with perfective aspect (The man pounded the nail).

Finally, although there was some priming found for instruments when the events were referenced as completed, the strength of this effect was weak in comparison to the results found in Ferretti et al. (in press). One plausible reason for this difference is that the perfect morpheme is more constraining toward the resultant states of events than the simple past form of verbs which references *entire* events as completed. That is, referencing entire events as completed may lead to stronger activation of common instruments because they are clearly more likely to be part of events than to be part of the states that follow the completion of the events.

Implications for semantic priming literature

The current results represent the first time in the semantic priming literature that verbs with aspectual morphemes have been used as primes in a short SOA priming task. Thus it is important to consider the implications of the current results to this literature. The most common explanation of word-word priming results is based on the notion of spreading activation (see Neely, 1991, for an extensive review). The original instantiations of semantic networks focused on noun representations (Collins & Loftus, 1975; Collins & Quillian, 1969). Further extensions implemented verb meaning into this framework (Gentner, 1975; Rumelhart & Levin, 1975). A verb's representation in Gentner's and Rumelhart and Levin's models included core meaning and thematic links to nodes that stood as placeholders for various thematic roles. However, it is important to note that these thematic links and nodes included minimal semantic content. For example, Rumelhart and Levin assumed that the agent link carried with it the information that any noun that could be placed in that node must be animate. In other words, the semantic content was restricted to general selectional restriction information. Because all locations and instruments were equivalent in terms of selectional restrictions, these models cannot account for the priming effects of Experiment 1.

If one were willing to expand current semantic networks, it could be assumed that a verb node becomes linked to noun nodes representing its common instruments and locations over time as a result of experience with situations to which the verb refers, and linguistic descriptions of them. These links could result from people noticing that, for example, arenas are associated with skating, or via word co-occurrence in speech and text. If all words or concepts that tend to co-occur frequently in situations are linked in

semantics, and/or orthography and phonology via an associative relation, then priming from verbs to typical role fillers might result. Note that the standard off-line method for determining whether two words or concepts share an associative link is word association norms. It has been argued in previous research that if subjects frequently produce one word given another in a word association task, than those concepts share an associative link, and facilitation should obtain in an on-line priming task (Neely, 1991). However, this type of explanation is problematic because *none* of the items in the current experiment was associated according to word association norms.[3]

Even more problematic for spreading activation theory, however, is explaining the different patterns of location priming found for the grammatical morphemes. If verbs were linked through undifferentiated links to their typical locations, one would have expected typical role fillers to be primed, regardless of the aspectual marking of the verb. The fact that location information was primed only with imperfect aspect suggests that this fundamental assumption of an undifferentiated spread of activation in semantic memory must be modified so that it can be constrained by syntactic information such as grammatical inflections on verbs. Thus, spreading activation networks could incorporate ad hoc assumptions well outside the scope of current versions to account for the present data, but they would no longer be in the same spirit upon which the notion of spreading activation was based.

---

[3]The word association norms are reported in Ferretti et al. (in press). In these norms, subjects were read the verbs and asked to produce the first word that came to mind. A verb-noun pair was considered associated if the noun was produced to the verb by greater than 5% of the subjects, which is a conservative criterion when compared to other priming experiments that have aimed to eliminate associatively-related items from their stimuli (e.g., Moss, Ostrin, Tyler, & Marslen-Wilson, 1995; Shelton & Martin, 1992).

These results are, however, consistent with Ferretti et al.'s claim that memory is structured so that when a verb is read or heard, situation-specific information via a situation schema is made available immediately. Moreover, the current experiment extends their results by showing how verb aspect constrains the activation of this information.

In conclusion, Experiment 1 investigated whether verb aspect can modulate the early activation of knowledge of common instruments and locations in situations when verbs are presented in isolation. When verbs were marked with imperfect aspect, knowledge of typical locations was activated. When the resultant states of the same situations were directly referenced, there was no evidence of activation of typical locations. Taken together the results indicate that aspect modulates the activation of knowledge about situations when people read verbs in isolation.

Chapter 3: Adjectival Participles and Situation Structure

Experiment 1 showed that when aspect references situations as ongoing, knowledge about situations is activated that is not available when the same situations are viewed as completed. These results are important because they demonstrate the close interplay between grammatical structure and event structure in lexical conceptual memory. Ideally, however, one would also want to show that this information is not simply made transiently available from verbs, but rather that it is used to drive crucial comprehension processes such as thematic role assignment. In this regard, it is important to recognize that the priming methodology employed in Experiment 1 provides no information about how people actively combine grammatical knowledge with event knowledge. Recall that in the priming task subjects silently read verb-phrase primes, and then named fillers into a microphone. There was some combinatorial processing involved with reading the verb phrases as subjects had to combine the verb with the aspect morpheme. However, when the target was displayed, the subjects were not required to combine or integrate the information made available by the verb phrase with the noun fillers they are naming. Thus, these results are limited to examining the transiently activated information from verbs. The main purpose of Chapter 3 was to investigate how people actively combine these kinds of information to constrain thematic role assignment during noun phrase interpretation.

A second purpose of the chapter was to extend the results to the agent and patient roles. These thematic roles are advantageous for the current purposes because human empirical studies and computational models suggest that participant information is salient in situations (Barsalou, 1988; Kolodner, 1983; Lancaster & Barsalou, 1997). It makes intuitive sense that

agents and patients, particularly human ones, are salient aspects of situations because people are social beings and generally attend to the activities of others. These roles also tend to be salient conceptually because agents and patients play well-defined roles in the causal structure of events. Agents tend to be associated with the initiating conditions because they *cause* and *control* events, while patients are typically associated with the resultant states because they tend to undergo a change of state (see Figure 2).[4]

Another important property of these roles is their relationship with the core arguments of transitive verbs. Agent and patient roles are by far the most frequent roles assigned to core arguments of transitive verbs (subject and object, respectively). Furthermore, the main clause (subject-verb-object) is the most frequent structure in English, and thus agent and patient role assignment is the most common form. In sum, the well-defined lexical/conceptual nature of transitive verbs makes them ideal for examining the influence of grammatical morphemes on constraining the activation of world knowledge of events.

A third, but equally important, purpose was to combine computational modeling with both on-line and off-line experimentation to gain further insight into how people combine morpho-syntactic information with knowledge of events during NP interpretation. The computational model makes the proposed interpretation process explicit, and is used to generate specific predictions about how people interpret the phrases on-line.

Chapter 3 begins by discussing the modifier-noun phrases that play a central role in the following research. Following this, the competition-integration model used to make predictions about the interpretation of these

[4]This should be taken as true of events in general. There are exceptions in which patients cause an event to occur because of their actions, and agents may also undergo a change of state as a result of the event.

phrases on-line is discussed. To provide a better understanding of how people assign thematic roles in these NPs, an off-line norming study is presented. These off-line norms provide motivation for setting of the parameters in the competition model. Next, the model's predictions are tested in two on-line experiments. Finally, Chapter 3 ends with a discussion of the implications of the results for theories of conceptual combination.

3.1 Adjectival Participles and Thematic Role Focus

To increase the ability of the grammatical morphemes to focus on the structural properties of events, the following experiments employed adjectival participles that modified head nouns (e.g., *arrested cop*). Adjectival participles are advantageous because they have properties that are consistent with both verbs and adjectives. For example, these participles inherit the argument structures of verbs from which they are derived. This suggests that participles derived from transitive verbs can be biased toward the thematic roles associated with their subject and object arguments. This bias is enhanced in adjectival participles because they have similar syntactic behavior as adjectives. Specifically, like adjectives in a modifying relation, only one argument of the adjectival participles is assigned a filler.

Whether or not an adjectival participle is biased toward focusing on the agent or patient role in events is determined by the type of grammatical inflection associated with the participle (see Haspelmath, 1994, for a review). *Present participles* such as "arresting" in the phrase *arresting cop* are biased toward the agent role in events, and *past participles* (*arrested cop*) are biased toward focusing on the patients in events (see Figure 4). Note that the terms "present participles" and "past participles" are used because

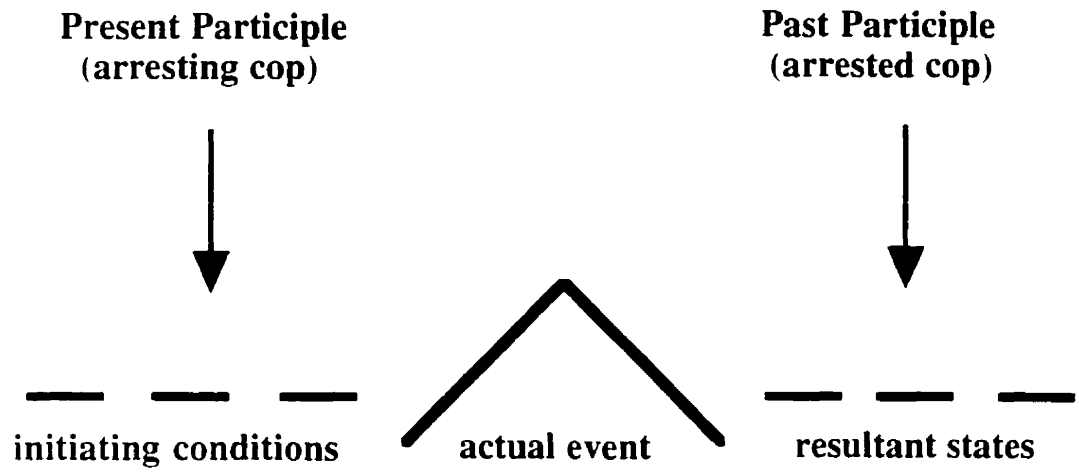

Figure 4. *Adjectival participles and their relation to the temporal and causal structure of events.*

they are the most frequent terms applied to these modifiers. However, present participles are sometimes referred to as *active* participles, and past participles are often referred to as *passive* or *resultative* participles. The approach taken here, following Haspelmath (1994), is that these two morphemes capture aspectual distinctions. The focusing properties of these two morphemes are derived from the fact that adjectives describing situations tend to be more time-stable (i.e., more stative) than the same situations described by verbs. For example, Haspelmath proposes that the only way in which "things" (i.e., head nouns) can be captured in such a time-stable manner is if the head noun is described in terms of *habits*, *results*, and *eventualities*. For example, in the following quote, Haspelmath explains why resultative participles are oriented toward the patient role:

> "Results of perfective telic situations are states that are naturally quite time-stable, and languages make extensive use of such resultative participles. These are primarily oriented toward the patient of the event because the patient undergoes a change as a result of the event and can therefore be easily characterized in terms of the event." (p. 164)

Thus, like perfect aspect, the resultative participle focuses on the patient role because it is the participant that is associated with the resultant states in events. This focusing property should hold for the past participles used in the following experiment because they were all derived from transitive verbs that involve participants that undergo changes of state.

Alternatively, the ability of present participles to define things in terms of habitual actions is what leads to the focus of these participles on the agent role. According to Haspelmath (1994), habitual situations are dynamic (non-stative) but time-stable to a sufficient degree that many languages, including English, use them as adjectives. Haspelmath further proposes that these types of participles (e.g., *traveling salesman*) are oriented toward

the agent "because habitual actions can be construed as properties of agents, not of patients like resulting states and eventualities" (p. 164). Finally, Haspelmath suggests that present participles can refer to current events (i.e., the salesman is traveling), although this is viewed as secondary to the habitual interpretation because these events are not viewed as time-stable to the same degree.

Taken together, the foregoing discussion suggests that past and present participles are biased toward highlighting the agent and patient roles in events, respectively. The following section discusses how these biases may combine with situation-specific knowledge to constrain noun phrase interpretation.

Thematic role assignment and NP interpretation: a proposal

The conceptual and syntactic properties of present and past adjectival participles make them ideal for studying the interactions between syntactic and conceptual knowledge of events during thematic role assignment. The constraint that only one of the participle's arguments are assigned a thematic role suggests that when people interpret these types of NPs there may be "competition" between two alternative thematic role assignments.[5] When the NPs have present participle modifiers, there will be a stronger bias toward assigning the head noun to the agent role and, conversely, when the participles are in their past form, there will be a strong bias toward the patient role. Note that if these syntactic biases combine with detailed knowledge about the events denoted by the participles then the focus created by the morphemes should be on event-specific agents and patients. Another source of information that contributes to this competition is how

[5] Depending on the event that is denoted by the participle, it is possible that other roles could compete for assignment, such as instrument (e.g., *sewing machine*) and location (*parking garage*) (see Gagné & Shoben, 1997).

typical the agents and patients are for the alternative roles of the participles. In Experiments 2, 3 and 4, each head noun varied in its typicality as a filler for *both* the agent and patient roles of its corresponding participle. When the head noun is typical for the role that is consistent with the focus of the participle, and atypical for the alternative role (*arresting cop*, *arrested crook*), there should be little competition between the competing thematic role assignments. Alternatively, when a head noun is atypical for the expected role, but typical for the alternative role (*arresting crook*, *arrested cop*), there should be a great deal of competition between the agent and patient role assignments. Of course, these are the two ends of the continuum. In reality, each head noun varies in how it fills the alternative roles of the participles, and thus there should be levels of competition that fall between these two extremes.

Finally, a number of researchers have claimed that a verb's argument structure includes only syntactically-relevant selectional restrictions such as *animacy* (Caplan, Hildebrandt, & Waters, 1994; Chomsky, 1965). Given that agents in events tend to be animate, whereas patients can be either animate or inanimate, the current study used transitive verbs that typically involved animate agents and patients and further, all head nouns were animate. Thus, in the current study, subjects had to use event-specific information that was more detailed than animacy to assign the head noun to their proper roles. For example, to assign *crook* to the patient role and *cop* to the agent role of *arrested crook* and *arresting cop*, respectively, subjects will have to use information that is specific to the types of entities that typically fill the alternative roles in the situations. This world knowledge could include information that cops have the authority under the law to arrest someone, or crooks tend to steal things (McRae et al., 1997).

The following research makes this proposal explicit by implementing a computational model of the competition process, and by utilizing both off-line and on-line experimental methodologies. This research provides important insight into the relationship between grammatical and conceptual knowledge about events.

3.2 A Competition-Integration Model

The main purpose of this section is to present a computational model that makes explicit the proposed combinatorial process, and to generate predictions about how people interpret these phrases on-line. This section begins by presenting the competition-integration model used to simulate the human response latencies. Following this, an off-line experiment is presented that provides insight into how people interpret these phrases off-line, and provides motivation for setting the free parameters in the model.

In the current competition model, information is combined according to a competition-integration algorithm that was developed originally by Michael Spivey and colleagues (McRae, Spivey-Knowlton, & Tanenhaus, 1998; Spivey & Tanenhaus, 1998; Spivey-Knowlton, 1996). The competition algorithm simulates the *process* of combining constraints on-line, and makes no claims about how the information constraints become represented. However, the model does assume that the constraints are represented in a manner that allows for them to be integrated immediately in parallel. Note that this does not imply that no theoretical stance is taken on the organization of the information in memory. On the contrary, the representational assumptions proposed by Ferretti et al. (in press) and McRae et al. (1997) are viewed as relevant.

In the current model (see Figure 5), the alternative thematic role interpretations (i.e., agent versus patient; large ovals in Figure 5) compete during processing until the activation of one interpretation reaches a criterion value. Because the sum of the activation of these interpretation nodes is always one, for one interpretation node to become more activated, the other interpretation node has to decrease its activation proportionally.

There are two informational constraints (rectangular boxes) that contribute to competition between the two interpretation nodes. The box on the left represents *Thematic Fit*, which captures each head noun's role/filler typicality rating for the agent *and* patient roles of their participle modifiers. Each head noun's agent rating supports the agent interpretation, and the patient rating for each head noun supports the patient interpretation.

The typicality of the head nouns for filling the alternative thematic roles of the participles used in Experiments 2, 3 and 4, were determined from a role/filler typicality norming study reported in McRae et al. (1997). The underlying assumption of this norming method that is crucial to the present research is that participle-noun phrases denote situations that happen in the world. Some of these situations tend to be common and others less common. Role/filler typicality norms index the plausibility of nouns as fillers for thematic roles of participles by measuring how typically the denoted situation occurs in the world. These norms differ from sentence ratings in that they abstract away from the linguistic structures and focus directly on knowledge about events or situations.

The validity of these norms as indexing thematic fit has been shown in previous research that demonstrated that they provide a better estimation than the standard methodology of rating entire phrases or sentences for

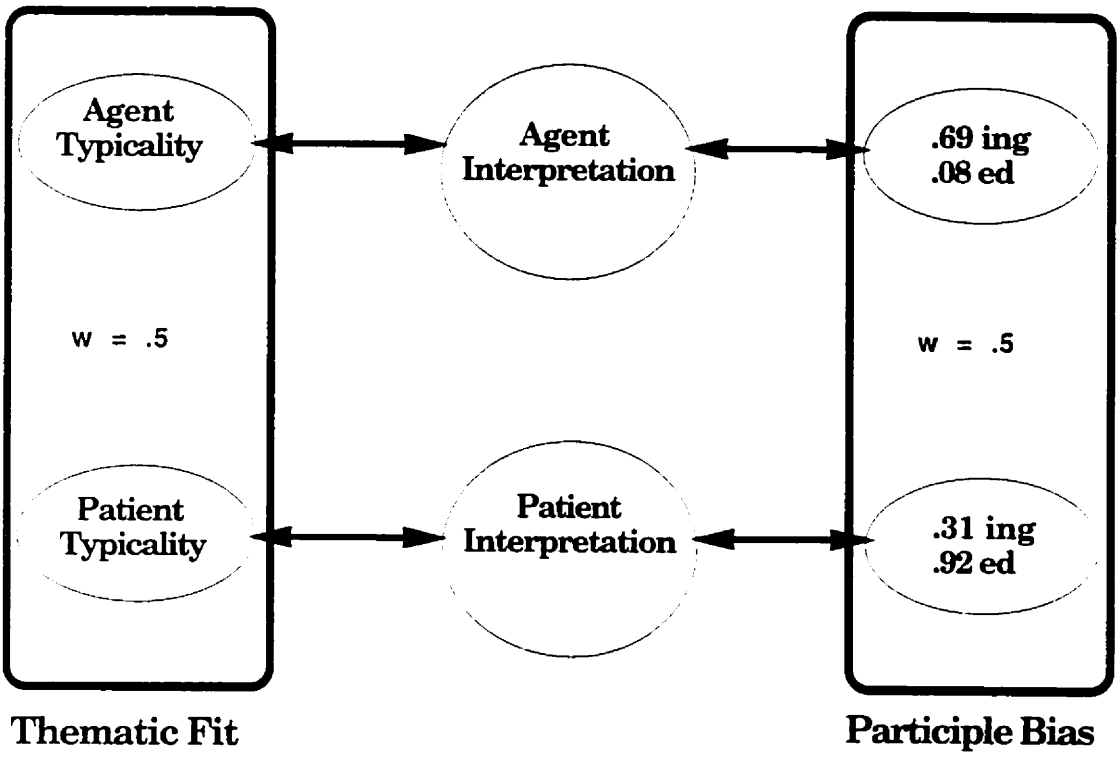

Figure 5. *A schematic diagram of the competition integration model.*

plausibility (Ferretti & McRae, 1999), and by research showing that the world knowledge tapped by these norms is computed and used immediately to constrain syntactic ambiguity resolution (McRae et al., 1998).

To estimate the typicality of the nouns in terms of filling the agent role of the participles, subjects rated agenthood on a 7-point scale as in (3.1), where 1 corresponded to very uncommon and 7 to very common.

(3.1) How common is it for a

crook ------

cop ------

guard ------

police ------

suspect ------

to arrest someone?

Subjects also rated how well each noun filled the patient role by answering questions similar to (3.2).

(3.2) How common is it for a

crook ------

cop ------

guard ------

police ------

suspect ------

to be arrested by someone?

In general, the NPs selected from this norming study were rated as a typical filler for one of the participle's roles, and relatively low for others. If the head nouns were most typical for the agent role of the participle they were classified as good-*agent NPs*, and if they were most typical of the patient role they were classified as good-*patient NPs*. Finally, although all

of the NPs used in Experiments 2, 3 and 4 were derived from this norming study, the three experiments contained slightly different, but highly overlapping sets of items from this norming study. The items used in each experiment and their corresponding role/filler typicality ratings are listed in Appendices B and C.

The second constraint that contributed to competition is *Participle Bias*, which captures the bias of present and past participles toward the agent and patient roles. The procedure used to estimate participle bias is discussed in detail below. The values of these biases sum to 1, and the values of the biases for the alternative morphemes were determined by fitting the off-line data, as described below.

A major assumption of the current model is that the amount of competition produced is proportional to the time it takes for subjects to indicate that a NP has a sensible interpretation. Competition in the model arises because thematic fit and participle bias simultaneously provide support for the agent and patient interpretation nodes. When the two constraints strongly support the same interpretation, the model will produce few cycles of competition. Alternatively, the model produces a greater number of cycles of competition when the constraints support both interpretations to a relatively equal degree. Constraint values that fall between these alternatives produce graded levels of competition. This attribute makes the competition produced by the model useful for capturing subtle interactions between the thematic fit of the head nouns and the biases of the present and past participles.

One Cycle of Competition in Model. The competition-integration algorithm produces competition using a three-stage process. In the first

stage, the normalized probabilistic support of each constraint for the two competing alternatives is computed according to Equation 1.

Equation 1 $S_{c,a}(\text{norm}) = S_{c,a} / \Sigma S_{c,a}$

In Equation 1, $S_{c,a}$ represents the activation of each of the constraint nodes, and is represented by the small ovals illustrated in Figure 5. The c represents the c^{th} constraint that is connected to the a^{th} interpretation node. $S_{c,a}(\text{norm})$ represents the normalized probabilistic support for the c^{th} constraint that is connected to the a^{th} interpretation node, and is computed by dividing the activation of $S_{c,a}$ by the sum of all of the $S_{c,a}$ activations within the same informational constraint.

In the second-stage, each normalized input is multiplied by a weight before feeding into its corresponding interpretation node (i.e., the large circles in Figure 5). Following this, the weighted inputs are summed to produce the interpretation node's activation.

Equation 2 $I_a = \Sigma [w_c * S_{c,a}(\text{norm})]$

In Equation 2, I_a represents the activation of interpretation node a, and w_c represents the weights shown in Figure 5. It is important to note that different sources of information are being combined simultaneously by this equation.

In the third stage, each interpretation node's activation is multiplied by its weighted inputs to provide feedback to each of the constraint nodes. Note that feedback to each of the constraint nodes is proportional to the

support it gave to the interpretation node. Following this, each constraint node's activation for the next cycle is calculated according to Equation 3.

Equation 3 $$S_{c,a(t+1)} = S_{c,a(t)}(\text{norm}) + I_a * w_c * S_{c,a(t)}(\text{norm})$$

In this equation, $S_{c,a(t+1)}$ represents the constraint node's activation at the next cycle, $S_{c,a(t)}(\text{norm})$ is that node's normalized activation at the previous cycle, and $I_a * w_c * S_{c,a(t)}(\text{norm})$ represents the feedback from the interpretation node that it supported during the previous cycle. Note that the different sources of information are indirectly resolving each other's ambiguities because I_a is the sum of their weighted inputs.

To begin the next cycle, $S_{c,a(t+1)}$ is normalized again to calculate its probabilistic support for interpretation node a. The net effect of this three-stage process (normalized recurrence) over time is that each constraint node accumulates normalized feedback. As a result of normalization, the constraint nodes that provided the strongest support for an interpretation become even stronger, and those that provided weak support begin to provide even less support. Eventually, one of the interpretation nodes accumulates sufficient probabilistic support to stop competition.

In the current model, competition begins with the information about participle bias and thematic fit contributing to competition. The model then produces competition until one of the interpretation nodes reaches a criterion value called the dynamic criterion, which is represented by Equation 4.

Equation 4 $$\text{Dynamic Criterion} = 1 - \Delta crit * cycle$$

In Equation 4, $\Delta crit$ is a constant that controls the rate of change, and cycle represents the number of cycles that have passed since competition started. According to this equation, the criterion that has to be met before competition stops decreases as the number of cycles increases (see Table 2). This equation captures the fact that comprehenders eventually make a judgement about the NPs. Finally, explorations with the model showed that a Δcrit of .007 provided reasonable behavior (e.g., was not producing ceiling effects in the activation of the interpretation nodes). This Δcrit is lower than used in previous models such as McRae et al. (1998) who used .008 to simulate self-paced reading, and Spivey and Tanenhaus (1998) who used .01 for eyetracking data. However, the lower value is appropriate in the current research because response latencies in the sense/nonsense task employed below tend to be longer than the reading times simulated with past models.

3.3 Estimating Participle Bias and Weights

There are two sets of parameters that need to be chosen; the participle biases, and the weights for Thematic Fit and Participle Bias. The following section discusses in detail how these biases were estimated.

Participle bias captures the support of the present and past participles toward the alternative thematic role interpretations. Given that previous research suggests that present and past participles are oriented toward agents and patients respectively (e.g., Haspelmath, 1994), the obvious choice for the values of these biases is full support of 1 for the interpretation node that matches the orientation of the participle (e.g., 1 to the patient interpretation node when a past participle is used), and no support for the interpretation node that does not match (0 to the agent interpretation node

Cycle	*Dynamic Criterion*
1	.993
2	.986
3	.979
4	.972
5	.965
10	.930
20	.860
30	.790
71	.503

Table 2

Change in the dynamic criterion as a function of cycles of competition.

for past participles). However, there are two possible complications: they are not categorically-biasing, and they are not equally biasing. Fortunately, there exists a method for setting parameters such as the participle bias and the weights for a model that is designed to predict on-line data.

The procedure adopted in this thesis is similar to that used by McRae et al. (1998). The weights in their model were motivated by an off-line sentence completion task. This task provided information about how frequently people made one of two competing interpretations (main clause, reduced relative clause) for their sentences. To set their weights, they found the combination of settings that produced activations in the alternative interpretations that were most similar to the frequency of the same interpretations in the off-line task. Setting weights in this manner provides a principled and theoretically-neutral approach (McRae et al., 1998).

To implement this procedure, an off-line task is needed that is equivalent to the gated completion norms used by McRae et al. (1998). Toward this end, Experiment 2 was conducted to gain insight into how people interpret this type of combination, at least in an off-line task.

3.4 Experiment 2

In Experiment 2, subjects provided a definition for participle-noun phrases. The two objectives of this experiment were to gain insight into how people interpret these types of NPs in an off-line setting in which there is no time pressure to respond, and to use the frequency with which people define the head nouns in the phrases as agents and patients to provide independent motivation for the parameters in the competition model. For instance, the proportion of phrases that are given agent and patient interpretation when the NPs have present and past participles should

provide insight into the bias of the different morphemes toward the agent and patient roles of the participles. If past and present participles are biased completely toward the patient and agent role, respectively, then both types of NPs should be interpreted as agents and patients about equally often.

Method

Subjects.

Forty-eight subjects participated for course credit. All subjects were native English-speaking psychology undergraduates from the University of Western Ontario.

Materials and Procedure.

The materials consisted of 2 booklets that contained 64 good-agent NPs, and 2 that contained 56 good-patient NPs. All of these items were selected from the previously reported norming study (McRae et al., 1997). The agenthood ratings of the 64 participle-agent pairs were significantly higher (M = 6.5, *range* = 5.7 - 6.9) than their patienthood ratings (M = 2.9, *range* = 1.3 - 6.6), $t_2(63) = 18.16$. In contrast, the patienthood ratings of the 56 participle-patient pairs were significantly higher (M = 6.0, range = 5.1 - 7.0) than their agenthood ratings, (M = 2.8, *range* = 1.0 - 6.6), $t_2(55) = 15.71$. The items and their corresponding role filler ratings are listed in Appendix B. Across the 2 sets of booklets, each participle appeared in their present and past forms. No participle or noun appeared more than once in each booklet. In addition, 4 other booklets were created that were identical to the 4 discussed above with the exception that the items appeared in reversed order.

The subjects were each given one version of the four booklets and were provided with the following instructions: "This experiment investigates

how people interpret noun phrases. In this study, you will be presented with many phrases. Your task is to read each phrase and think of a likely meaning for it. Pretend that you just heard the phrase in a conversation. What would be the meaning of the phrase that seems most natural to you? For each phrase, write the meaning that seems most natural to you. Please avoid vague interpretations. For some of the phrases, it may be difficult to come up with a meaning. Just do the best that you can". Note these instructions are similar to those used to examine how people define noun phrases off-line (e.g., Wisniewski, 1996). Each booklet took approximately 30 minutes to complete.

Design.

Percentage of patient interpretations for each condition were analyzed by two-way analyses of variance. The factors of interest were grammatical marking (present vs. past participle), which was within items (F_2), and noun phrase type (good-agent NP vs. good-patient NP) which was between items.

Results

Each item was coded for whether they were given "agent", "patient", or "other" interpretations in their definitions. A definition was coded as an agent interpretation if it described the head noun as being the agent in the event denoted by the participle. For example, the definition "a cop that arrested someone", which was frequently provided for *arresting cop,* was coded as an agent interpretation because the head noun was defined as the agent in the event denoted by the participle. Alternatively, definitions such as "a crook who was arrested" for *arrested crook* were scored as patient interpretations because the head noun in the phrase was defined as a patient in the event denoted by the participle. The "other" category was

created because some responses clearly did not fit either the agent or patient categories. These types of phrases were usually a result of subjects not providing much detail in their definitions (e.g., subjects gave one word descriptions that were associates of either the participle or head noun rather than a definition of what the phrase meant, such as writing "school" as a definition for *graded student*), or because they defined the head noun in the phrases neither as an agent or patient in the event (e.g., one subject provided the definition "army at war" for *attacked infantry*). Sometimes subjects implied in their definitions that the head noun was an agent or patient for a specific event without being explicit. In these cases, intuition was used to code the interpretation as either "agent", "patient" or "other". For example, one subject wrote "frightened and took off" as a definition for *startled bird*. This definition was coded as a patient interpretation because although the head noun was not included in the definition, it strongly implied that the bird was frightened and took off. Such examples were coded as agent or patient interpretations only when these interpretations were strongly implied, otherwise they were coded as "other".

It should be noted that 90% of definitions were either agent and patient interpretations. Due to the low number of "other" interpretations, they were removed from the analyses, and the proportion of agent and patient interpretations were normalized so they summed to 1.[6] Thus the analyses were performed only on those phrases in which subjects clearly categorized or strongly implied that the head noun was either an agent or patient in the event denoted by the participles. The normalized values are presented in Table 3.

---

[6]Note that the analyses were also performed with the "other" category included, but they did not differ in any significant way from the results reported below.

	Interpretation Type			
	Agent		Patient	
Phrase Type	*M*	*SE*	*M*	*SE*
Good-Agent				
(-ing)	.91	(.02)	.09	(.02)
(-ed)	.17	(.02)	.83	(.02)
Good-Patient				
(-ing)	.32	(.03)	.68	(.03)
(-ed)	.03	(.02)	.97	(.02)

Table 3

Proportion of head nouns interpreted as agents and patients in Experiment 2.

Visual inspection of the data in Table 3 indicates that for good-agent NPs subjects gave considerably more agent interpretations than patient interpretations when the NPs had present participles. Alternatively, the same NPs with past participles were given patient interpretations more often than agent interpretations. Interestingly, the means for the good-patient NPs indicate that subjects gave patient interpretations more often than agent interpretations for phrases that contained present *and* past participles. Thus, for agent NPs the different morphemes are strongly biasing the subjects toward the alternative roles, whereas for patient NPs only the past participle is highly constraining toward a single role.

The results of the ANOVA are consistent with these observations. Specifically, the two-way interaction between grammatical marking and NP type was significant, $F2(1, 118) = 83.99$. The large percentage of patient interpretations made for NPs with both types of participles lead to a significant main effect of NP type, indicating that overall subjects made significantly more patient interpretations for good-patient NPs ($M = .83$, $SE = .02$) than for good-agent NPs ($M = .46$, $SE = .04$), $F2(1, 118) = 451$. Finally, there was a main effect of morpheme type indicating that subjects made more patient interpretations for NPs with past participles ($M = .90$, $SE = .02$) than for those with present participles ($M = .37$, $SE = .03$), $F2(1, 118) = 215$.

Discussion

When the NPs contained past participles, subjects consistently gave patient interpretations regardless of whether the head nouns were good agents or patients for the participles. This result supports the notion that the past participle morpheme focuses on the participant that is part of the resultant states of events (Haspelmath, 1994), and that this focus is highly constraining toward a single role. In contrast, the overall bias of the

present participle was not as constraining toward the agent role as only 64% of the subjects gave agent interpretations.

This latter result was unexpected, as the present participle was anticipated to be more constraining toward the agent role. One account that provides a good explanation of these biases is Haspelmath's suggestion that the less constraining nature of the present participles is due to them not being as time-stable as past participles which focus on *states* that result from an event. This property of present participles provides a natural explanation for why subjects generated a number of patient interpretations when the head nouns were good patient fillers but not when they are good agents. Because the morpheme is heavily biased toward the agent role subjects will almost always make the favored agent interpretation. However, as the head noun becomes a better patient and a worse agent, the strong bias can be overcome and patient interpretation is made by interpreting the participle as a finite verb. That is, subjects can interpret the participle as a verb and the head noun as a direct object (i.e., as a verb phrase).

The definitions for *arresting crook* NPs provide support for this claim by indicating that subjects made these interpretations by generating events in which the head nouns filled the roles for which they are most typical (e.g., "a cop arresting a crook", or "a crook arrested by someone"). Note that linguistically, to make these interpretations from the NPs, subjects must be interpreting the participles as finite verbs rather than as non-finite participles. In other words, they are reconfiguring the syntactic structure of the NPs.

There are, however, other possible reasons why subjects may be more willing to make VP than NP interpretations for good-patient NPs with

present participles. The first is that McRae et al. (1997) showed that the agent roles for the transitive verbs used in their study tended to be more "well-defined" than the patient roles; that is, the verbs generally admit a wider range of fillers for the patient than for the agent role. For example, just about any person can be arrested (including cops), whereas a limited number of people can arrest others (not likely a crook). Thus, if patient roles for the items used in the current experiment generally admitted a wider range of fillers, then for *arresting crook* pairs it may be harder to produce an agent interpretation. This possibility could account for the finite-verb patient interpretations. Alternatively, for *arrested cop* pairs, a plausible patient interpretation consistent with the syntactic structure of the NP is easy to generate and thus subjects do not interpret the modifier as VPs as frequently. This possibility is investigated in Experiment 3.

Another possible contribution to this effect could be that the most common structure in English is the main clause which has a subject-verb-object form, which thematically is typically an agent-verb-patient structure. Thus, it is possible that people's sensitivity to word order constraints in English could be contributing to the ease of generating verb phrases for *arresting crook* NPs. That is, the fact that a noun is usually a patient when it directly follows a verb in English might be contributing to subjects' tendency to produce patient interpretations. Agents can follow verbs in English in passive constructions ("a man was arrested by the cop"), but the frequency of these constructions in comparison to main clauses is very low. It would be interesting to examine whether the same results would be obtained in languages that do not have strong word order constraints (e.g., Dutch and German).

3.5 Simulating Off-Line Interpretations

One objective of Experiment 2 was to provide independent motivation for setting the Participle Biases and the weights for both the Thematic Fit and Participle Bias. The proportion of subjects that gave agent and patient interpretations when the NPs were marked with present and past participles provide insight into the setting of the free parameters in the model. The purpose of simulating Experiment 2 was to estimate the parameter settings that best captured the proportion of agent and patient interpretations made by subjects for the same NPs in the off-line task. Once these parameters are set, the model will then be used to make predictions about the on-line performance of subjects in Experiments 3 and 4.

Each of the NPs employed in Experiment 2 were entered separately into the model, and a number of simulations with different parameter settings were performed. The model with the settings that provided the best fit to the results of Experiment 2 was retained. For each NP, the model began with the information about participle bias and thematic fit contributing to competition. For example, to simulate *governing villager*, the head noun's agent rating (3.2) and patient rating (6.5) were first entered into the model along with the participle bias associated with present participles, which was a free parameter in these simulations. In the first simulation, this bias was set to full support for the agent interpretation node and thus the patient interpretation received no support. Further, the weights in the model for Thematic Fit and Participle bias were originally set so that they were equal. The model then produced competition using the three step normalized recurrence procedure discussed above. When the model reached criterion, competition ended and the activation of the alternative interpretation nodes were recorded. For *governing villager*,

when the dynamic criterion was reached the patient interpretation node was strongly activated (.84) in comparison to the agent interpretation node (.16) (see Appendix C for a complete example of the model processing this NP over time). Thus the model made a patient interpretation and, further, the activation of the patient interpretation was considered to correspond to the prediction that 84% of subjects would make a patient interpretation for this NP. This process was repeated for every item in Experiment 2. The mean proportion of responses predicted by the best fitting model for present and past participles by filler type (good-agent NP, good-patient NP) are illustrated in Figure 6.[7] As illustrated in Figure 6, the activation of the model's agent and patient interpretation nodes closely capture the proportion of subjects that gave agent and patient interpretations for the same NPs in Experiment 2.

The parameter settings that provided the best estimate of the off-line interpretation frequencies were weights of .5 and .5 for participle bias and thematic fit, and biases of .69 and .31 toward the agent and patient interpretations, respectively, when simulating present participles. Alternatively, when simulating past participles, the bias toward the agent interpretation was .08 and the patient interpretation received a bias of .92.

The setting of parameters that provided the best fit to the results in Experiment 2 have implications for the proposed combinatorial process. First, the equal weight settings for thematic fit and participle bias suggests that world knowledge of situations and the syntactic information associated with the participles are weighted equally by subjects when they interpret the phrases. Second, the setting of the participle biases for the morphemes

[7]The activation of agent role for the same conditions can easily be calculated by subtracting the patient activation from 1.

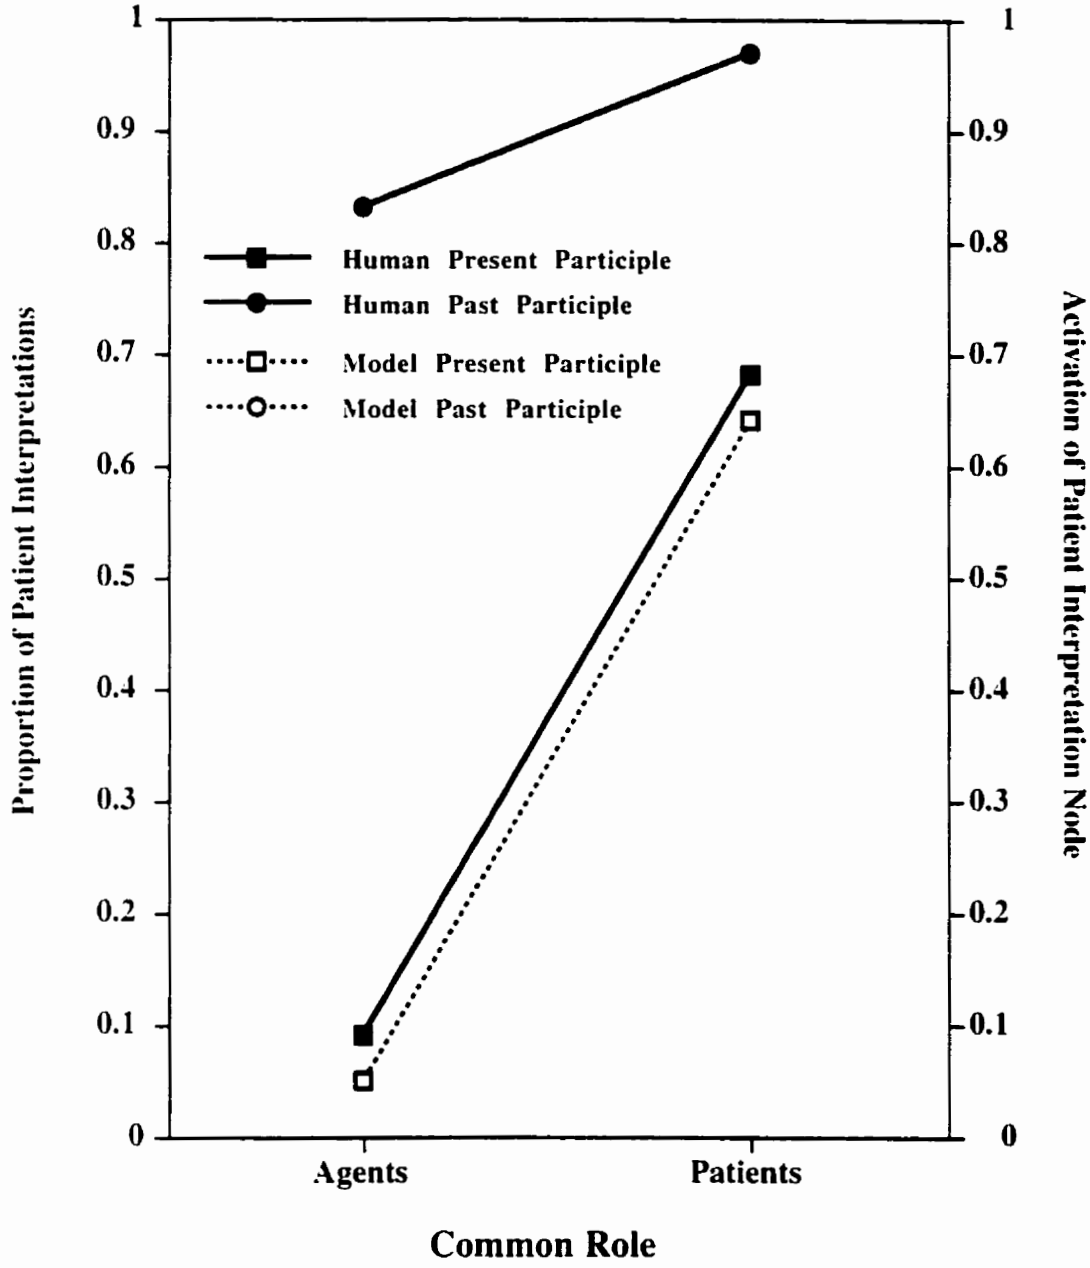

Figure 6. *Using activation of the patient interpretation node to simulate proportion of patient interpretations in Experiment 2. Note that the model's data for present participles (-ing) are difficult to see because they overlap so highly with the human data.*

suggests that present participles are not as biased toward a single role as are past participles. These points will be revisited below.

3.6 Predicting On-line Modifier-Noun Interpretation

The goal of this section was to use the parameter settings obtained from the model that provided that best simulation of Experiment 2 to make response latency predictions about how subjects would interpret these types of phrases on-line in Experiment 3. The 58 good-agent and 53 good-patient NPs that contributed to the results of Experiment 3 were used in a new simulation. Similar to the procedure used to simulate the NPs in Experiment 2, each NP and its corresponding role/filler typicality ratings was entered into the model. The number of cycles of competition produced by the model to reach criterion for each NP was recorded along with the activations of the interpretation nodes.

The means for each type of NP are illustrated in Figure 7. This simulation represents the model's prediction about the behavior of the humans in the on-line task. There are two key aspects of the model's output that are important to note, including the mean number of cycles of competition produced for the different types of NPs, and the mean activation of the interpretation nodes for the same NPs. The present discussion focuses on the response latency predictions, and thus the mean cycles of competition produced by the model is of particular interest. However, the model's interpretations will be addressed in detail later.

The model makes specific predictions about how quickly subjects should interpret the phrases. For example, the model produced more competition for phrases with mismatching morphemes and head nouns (*arrested cop*, *arresting crook*) than for phrases in which the two types of

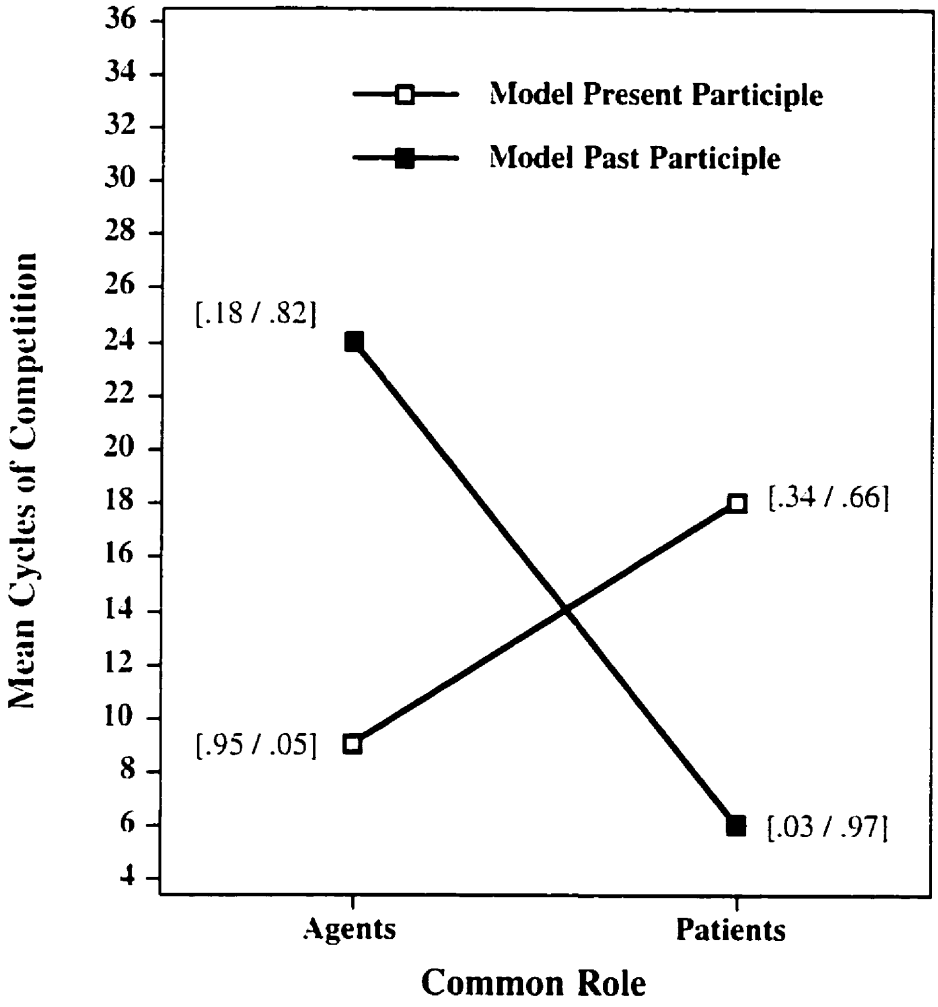

Figure 7. *The model's prediction for response latencies in Experiment 3. The values in brackets indicate the mean activation of the agent and patient interpretation nodes for each condition.*

information are consistent (*arresting cop*, *arrested crook*). Thus the model predicts subjects should take longer to respond to NPs that have mismatching sources of information. Another important aspect of the model's pattern of competition was the relative difference in cycles (9 cycles) for the two types of NPs with present participles (*arresting cop*, *arresting crook*) versus the difference in competition for the NPs with past participles (*arrested cop*, *arrested crook*) (18 cycles). Thus the model is predicting that the differences in response latencies should be less for NPs with present participles than for NPs with past participles.

To ground the model's prediction in inferential statistics, the mean cycles of competition were analyzed by two-way analyses of variance. The factors of interest were grammatical marking (present vs. past), which was within items, and NP type (good-agent vs. good-patient) which was between items. Grammatical marking and NP type interacted, $F_2(1, 109) = 136.5$. Thus the model predicted that the same interaction should be found in ANOVA for response latencies. Planned comparisons for good-patient NPs revealed that the model produced fewer cycles of competition when the NPs were in their past versus present form, $F_2(1, 109) = 53.0$. Alternatively, planned comparisons for good-agent NPs revealed that the model produced fewer cycles of competition when the NPs were in their present than past form, $F_2(1, 109) = 91.0$. Thus, the model predicted that subjects in Experiment 3 should be significantly faster to respond when the expected filler is consistent with the bias of the participles (*arresting cop*, *arrested crook*) than when they are inconsistent (*arrested cop*, *arresting crook*).

3.7 Experiment 3

Experiment 3 examined how people combine their knowledge of events with morpho-syntactic information during on-line NP interpretation. I used a task commonly employed in the conceptual combination literature called a sense-nonsense task (e.g., Gagné & Shoben, 1997; Murphy, 1990). In the present version, subjects indicated as quickly and accurately as possible whether modifier-noun NPs presented in isolation on a computer screen had sensible interpretations. It is standardly assumed that deciding whether or not a modifier-noun phrase has a sensible interpretation involves combining the meaning of the constituent concepts. The current sense-nonsense task should therefore provide insight into how people combine event-specific knowledge with grammatical inflections to constrain thematic role assignment, and how people use event knowledge to combine concepts in modifier-noun phrases in general.

The predictions were as follows: subjects should interpret good-agent NPs more quickly when they are in their present (*arresting cop*) than past form (*arrested cop*), because the present participle focuses on the agent in events. Thus, the focusing properties of these morphemes should combine with event knowledge to restrict the modifying relation to the agent role, and this process should lead to facilitation in responses and more frequent "yes" judgements when the head nouns are good agents. Alternatively, good-patient NPs should be interpreted more quickly and judged more sensible when they are in their past (*arrested crook*) than present form (*arresting crook*). This should occur because the past participle should combine with event knowledge to create a modifying relation that is highly constraining toward patients in specific events.

To ensure that any interaction between the morphemes and thematic fit of the head noun toward the alternative roles is due to event-specific knowledge, and thus not simply due to head nouns that are just more likely to be agents or patients across events in general, an unrelated control condition was created in which the head nouns were repaired with other participles to form unrelated NPs that still had sensible interpretations (e.g., *kicked king*). If the head nouns for the good-agent NPs are biased toward being good agents in general, then the unrelated NPs with the same head nouns should be also interpreted more quickly when they are modified by present participles. Moreover, if the head nouns for the good-patient NPs are more patient-like in general for most events, then subjects should interpret the unrelated NPs also more quickly when they are in their past than present form. Finally, if event-specific knowledge is constraining thematic role assignment, an interaction between morpheme and NP type should be obtained for the related NPs only.

Method

Subjects

Eighty subjects participated for course credit, 10 per list. All subjects were native English-speaking psychology undergraduates from the University of Western Ontario who had normal or corrected-to-normal visual acuity.

Materials

Seventy-two NPs were created by pairing participles with a good-agent head noun, and another 60 NPs were formed by pairing participles with a good-patient head noun. The typicality of the head noun for the agent and patient roles of each participle was determined from the role/filler typicality

study presented earlier. The typicality values for the NPs used in the current study are listed in Appendix D.

The agenthood ratings of the 72 good-agent NPs were significantly higher (M = 6.5, *range* = 5.7 - 6.9) than their patienthood ratings (M = 3.0, *range* = 1.3 - 6.6), $t_2(57) = 17.13$, $p < .001$. In contrast, the patienthood ratings of the 60 good-patient NPs were significantly higher (M = 6.1, range = 5.1 - 7.0.) than their agenthood ratings, (M = 2.7, *range* = 1.0 - 6.6), $t_2(52) = 16.07$, $p < .001$.

Eight lists were constructed that included 4 with good-agent NPs (*arresting cop*), and 4 consisting of good-patient NPs (*arrested crook*). Each of the lists with good-agent NPs had 18 related NPs with present participles (*arresting cop*), 18 related NPs with past participles (*arrested crook*), 18 unrelated NPs with present participles (*punishing nurse*), 18 unrelated NPs with past participles (*kicked king*). Note that the unrelated trials were formed by repairing the typical head nouns with participles so that they were now plausible but not typical fillers. Each list also included the same 72 filler trials which consisted of nonsensical participle-noun combinations (e.g., *brewed ant*). Each of the 4 lists that contained good-patient NPs included 15 related NPs with present participles (*adopting baby*), 15 related NPs with past participles (*adopted baby*), 15 unrelated NPs with present participles (*attacking audience*), 15 unrelated NPs with past participles (*attacked audience*), and 60 filler trials that consisted of nonsensical participle-noun combinations (*brewed ant*). For both NP types, each participle appeared in their present (*arresting*) and past form (*arrested*), and modified both a related (*arresting cop*) and unrelated, but plausible (*loving cop*), head noun. Finally, 20 practice trials were created by pairing participles with a head noun. Half of these NPs had sensible interpretations

and the other half were nonsensical. Moreover, half of the practice trials were in their present form and half were in their past form.

Procedure

Stimuli were displayed on a 14-inch Sony Trinitron monitor controlled by a Macintosh LCIII using PsyScope (Cohen et al., 1993). Each trial consisted of the following events: first the word "Ready?" was presented in the center of the screen until the subject pushed a button indicating they were ready; 500 ms following this a participle-noun combination was presented until the subject responded. The intertrial interval was 1500 ms. Subjects were instructed to read the two-word phrase in the order that the words were given, and to decide as quickly and accurately as possible whether the phrase had a sensible interpretation. The buttons used to indicate "Yes" and "No" responses were randomized across subjects to control for any influence of handedess on response latencies.

Response latencies were recorded with millisecond accuracy via a CMU button box that measured the time between the onset of the noun phrases and the button press. Testing sessions began with the practice trials. It took approximately 20 minutes to complete the task.

Design

Response latencies and percent sensible responses were analyzed by four-way analyses of variance. The factors of interest were grammatical marking (present vs. past), which was within subjects and items, relatedness (related vs. unrelated events) which was also within subjects and items, and NP type (good-agent vs. good-patient) which was between subjects and items. Note that a trial was excluded from the latency analyses if the response indicated it was nonsensical. List and rotation group was used as a between item variable to stabilize variance due to rotating subjects

and items across different lists. Finally, after Experiment 3 had been conducted, it was discovered that some NPs were mistakenly marked with the past tense form of the verb (drove), rather than the past participle (driven) form. These items were removed from the analyses, which then were based on 58 good-agent NPs and 53 good-patient NPs. Note that this remains a large number of items for an experiment of this type.

Results

Analyses of Variance

Response latencies. Latencies greater than three standard deviations above the grand mean were replaced by that value (1% of the scores). Response latencies are presented in Table 4. The three-way interaction among grammatical marking, relatedness, and NP type was significant, $F1(1, 72) = 9.32$, $F2(1, 103) = 4.00$. This interaction occurred because grammatical marking and NP type interacted for related NPs, $F1(1, 72) = 12.56$, $F2(1, 103) = 14.68$, but not for unrelated NPs, both Fs < 1. Planned comparisons for related items revealed that good-patient NPs were responded to more quickly when they were in past than present form, $F1(1, 72) = 12.33$, $F2(1, 103) = 4.62$. Alternatively, good-agent NPs were responded to more quickly in their present than past form, $F1(1, 72) = 5.54$, $F2(1, 103) = 7.08$. In contrast, planned comparisons for unrelated items revealed that the participle patient items were not responded to differently when they were in their past versus present form, $F1 < 1$, $F2 < 1$, and that good-agent NPs also were not responded to differently when they were in their past versus present forms, $F1 < 1$, $F2 < 1$.

Grammatical marking and role interacted because NPs with present participles were responded to more quickly than those with past participles for good-agent NPs, whereas the opposite was true for good-patient NPs,

	Related Trials		Unrelated Trials	
Phrase Type	*M*	*SE*	*M*	*SE*
Good-Agent				
(-ing)	1805	49	2117	58
(-ed)	2005	64	2175	90
	200**		58	
Good-Patient				
(-ing)	1934	55	2124	72
(-ed)	1765	48	2151	58
	169**		27	

**Significant by subjects and items

Table 4

Mean response latencies (ms) for good-agent and good-patient NPs in Experiment 3.

$F_1(1, 72) = 5.94$, $F_2(1, 103) = 5.98$. The interaction between relatedness and role, and the interaction between relatedness and marking were not significant, all F's < 1. There was a main effect of relatedness; related items were responded to more quickly (M = 1904 ms, SE = 45 ms) than unrelated items (M = 2156 ms, SE = 53 ms), $F_1(1, 72) = 51.98$, $F_2(1, 103) = 32.73$. The main effect of role and the main effect of marking were not significant, all F's < 1.

Finally, it should be noted that overall subjects responded relatively slowly when making "nonsense" judgements (M = 2380 ms, SE = 18 ms).

Percent sensible responses. Percentage of sensible responses are presented in Table 5. The three-way interaction among grammatical marking, relatedness, and role was significant, $F_1(1, 72) = 10.53$, $F_2(1, 103) = 6.33$. This interaction occurred because grammatical marking and NP type interacted for related NPs, $F_1(1, 72) = 19.07$, $F_2(1, 103) = 25.71$, but not for unrelated NPs, both F's < 1. Planned comparisons for related items revealed that good-patient NPs were judged as sensible more often when they were in their past than present form, $F_1(1, 72) = 9.95$, $F_2(1, 103) = 8.08$. Alternatively, good-agent NPs were judged as sensible more often in their present than past form, $F_1(1, 72) = 10.94$, $F_2(1, 103) = 11.19$. Planned comparisons for the unrelated items revealed that the good-patient NPs were not judged as sensible more often when they were in their past versus when they were in their present form, $F_1 < 1$, $F_2 < 1$. Moreover, good-agent NPs also were not responded to as sensible more often when they were in their past versus present forms, $F_1 < 1$, $F_2 < 1$.

Grammatical marking and role interacted because NPs with present participles were judged as sensible more often than those with past participles for good-agent NPs, whereas the opposite was true for good-

	Related Trials		Unrelated Trials	
Phrase Type	*M*	*SE*	*M*	*SE*
Good-Agent				
(-ing)	94	1.3	74	2.9
(-ed)	85	2.0	76	2.9
	9**		2	
Good-Patient				
(-ing)	84	1.7	69	2.6
(-ed)	92	1.6	73	2.8
	8**		4	

**Significant by subjects and items

Table 5

Percentage of good-agent and good-patient NPs judged sensible in Experiment 3.

patient NPs, $F1(1, 72) = 7.34$, $F2(1, 103) = 12.42$. The interaction between relatedness and role, and the interaction between relatedness and marking were not significant, all F's < 1.2. Finally, there was a main effect of relatedness; related items were responded to as sensible more often (M = 89%, SE = 1%) than unrelated items (M = 73%, SE = 1%), $F1(1, 72) = 54.91$, $F2(1, 103) = 87.10$. Overall, good-agent and good-patient NPs were responded to as sensible equally often, $F1 < 1$, $F2 < 1$, and NPs with present participles were responded to as sensible as often as those with past participles, $F1 < 1$, $F2 < 1$.

Correlation Analyses

Correlational analyses were conducted to complement the analyses of variance by investigating how the agent and patient ratings for the head nouns correlate with response latencies for the related items. Although this was an ANOVA design, there was variation in the role/filler typicality ratings; that is, they vary on a relatively truncated continuum. The patient ratings for the 53 head nouns of the good-patient NPs ranged from 5.1 to 7.0, and the agent ratings for the 58 head nouns of the good-agent NPs ranged from 5.7 to 6.9. If subjects are sensitive to this continuum, then for example, response latencies and percent sensibility judgements for good-agent NPs such as *arresting cop* should correlate with agenthood rating.

Furthermore, the variation in ratings for the head nouns for their atypical role should be informative about the sensitivity of subjects to this factor. For example, the agent ratings of the head nouns used in the good-patient NPs ranged from 1.0 to 6.6. Thus, although the overall mean agent ratings for these head nouns were low (2.7), there was variation to which subjects may be sensitive. This pattern was also true for the good-agent NPs which had head nouns that ranged between 1.3 and 6.5 in patient typicality.

Phrase Type	*Agent Rating*	*Patient Rating*
Good-Agent		
(-ing)	-.37**	-.03
(-ed)	-.13	-.11
Good-Patient		
(-ing)	.01	-.30**
(-ed)	-.07	-.30**

**indicates a significant correlation

Table 6

Correlations between patient and agent ratings and response latencies for all NPs in Experiment 3.

Phrase Type	*Agent Rating*	*Patient Rating*
Good-Agent		
(-ing)	.06	.02
(-ed)	-.22*	.27**
Good-Patient		
(-ing)	.36**	.08
(-ed)	-.07	.24*

**indicates a significant correlation
*indicates a marginal correlation

Table 7

Correlations between patient and agent ratings and percentage of sensible responses for all NPs in Experiment 3.

The results of the correlational analysis for response latency and percentage of sensible responses are discussed in turn. The correlations between role/filler typicality rating and response latency are illustrated in Table 6, and Table 7 illustrates the same correlations with percentage of sensible responses.

Response Latencies. There was a significant negative correlation between agent rating and response latencies for good-agent NPs with present participles (*arresting cop*), demonstrating that subjects responded more quickly as the agent rating of the head noun increased. There was also a significant negative correlation between patient rating and response latencies for good-patient NPs with past participles (*arrested crook*). These correlations show that subjects were sensitive to differences in the thematic fit for head nouns of NPs in which the morphemes and head nouns were consistent. The correlation analyses also indicated that the variation in ratings for the same head nouns for the alternative roles of the participles for these NPs did not have a significant influence. That is, subjects' performance on this task was not influenced by thematic fit for the agent role of phrases like *arrested crook*, nor for the patient role for phrases like *arresting cop*.

When the fillers and head nouns were inconsistent (*arrested cop*, *arresting crook*), subjects' performance was influenced by thematic fit only for the *arresting crook* NPs. For these NPs, there was a significant negative correlation between patient rating and response latencies. Thus subjects responded more quickly as the head nouns increased in patient rating. Although this result is counter-intuitive, it is entirely consistent with the model's prediction that subjects will make a number of patient

interpretations for these types of NPs. This result is explored further in the section testing the model's predictions.

Percent sensible responses. The correlations for the percentage of sensible responses were fairly consistent with those for the response latency data. Specifically, there was a positive correlation between percentage of sensible interpretations and patienthood rating for good-patient NPs with past participles (*arrested crook*), indicating that subjects judged more phrases as sensible as the patient rating of the head noun increased. For good-agent NPs with present participles (*arresting cop*), agent ratings did not correlate with percent sensible responses, although this was due a ceiling effect. That is, subjects almost always gave these NPs sensible responses.

When the head nouns and participles were inconsistent (*arrested cop*, *arresting crook*), the role filler typicality ratings again influenced the percentage of responses. For example, the patienthood ratings for the head nouns for the good-agent NPs were positively correlated with percentage of sensible interpretations when the participles were in their past form (*arrested cop*). Interestingly, there also was a marginal negative correlation for the same items in terms of agenthood ratings. Thus subjects found these NPs more sensible as their head nouns increased in patienthood rating, and less sensible as the head nouns increased in agenthood rating. For good-patient NPs that mismatched with participle morpheme (*arresting crook*), there was a significant positive correlation between percent sense and agent rating. Surprisingly, there was no correlation between patienthood rating and percent sense for these items. This suggests that although subjects are faster to judge these NPs as

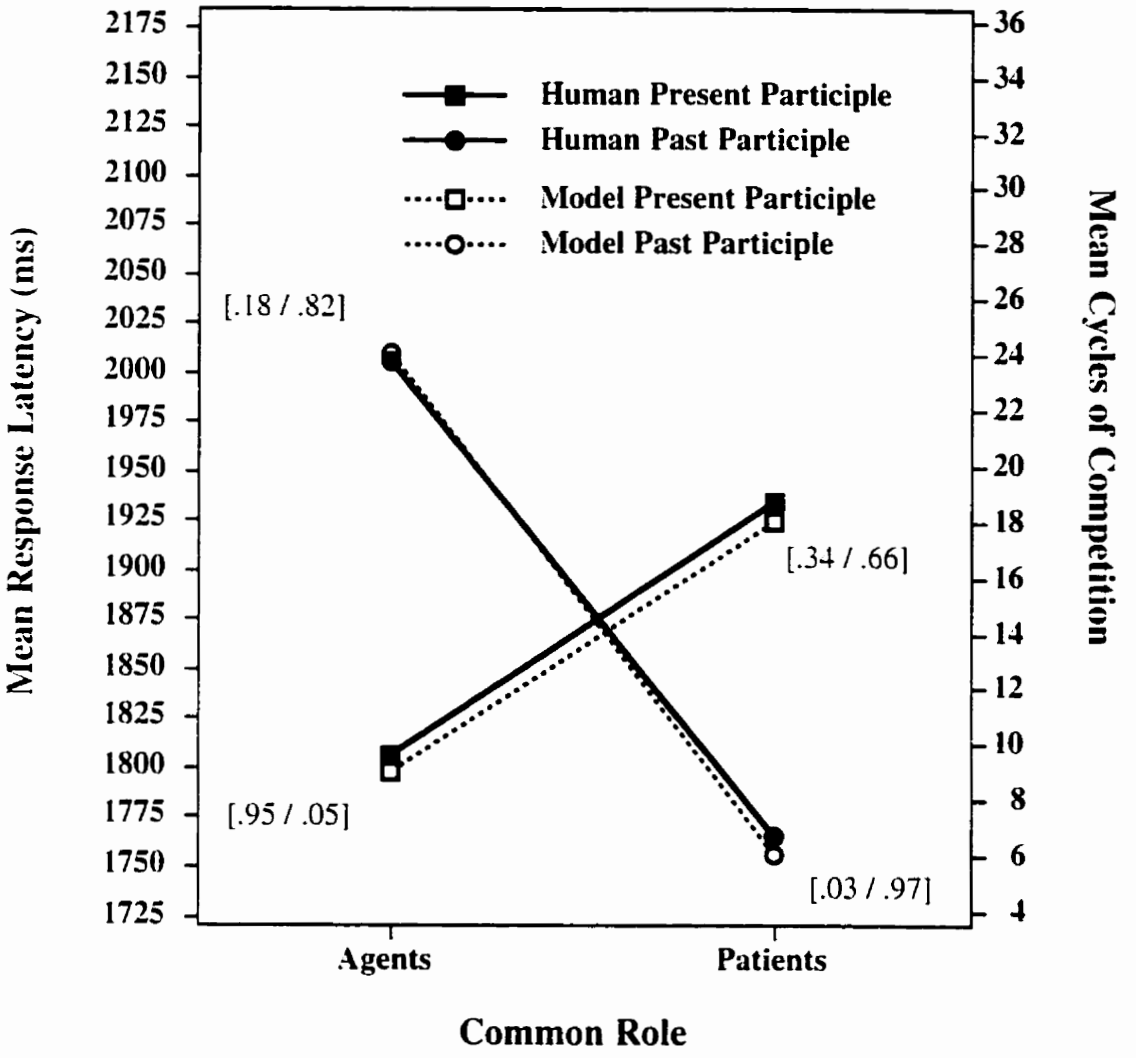

Figure 8. *Experiment 3 response latencies and the model's corresponding predictions. The values in brackets indicate the mean activation of the agent and patient interpretation nodes for each condition.*

sensible as the head nouns increased in patienthood rating, they do not judge them as sensible more often.

Model Predictions

The goal of the following section was to examine how well the model's predictions (i.e., the two-way interaction between grammatical marking and NP type) about response latency matched the obtained values. Figure 8 presents the model's predictions along with the corresponding human response latencies. As illustrated in Figure 8, the model captures all major aspects of the response latency data. Specifically, when past participles were used the model correctly predicted how difficult the phrases were to interpret: when the head nouns were typical agents (*arrested cop*) the model made the correct prediction that subjects should take much longer to respond; when the head nouns were typical patients (*arrested crook*) the model also correctly predicted that the subjects should have responded more quickly. Alternatively, when present participles were used, the model again predicts the relative difficulty that subjects had interpreting the phrases with good agent (*arresting cop*) and patient (*arresting crook*) head nouns.

So far the discussion has focused on the model's response latency predictions and, hence, the model's cycles of competition for each of the conditions have been of most interest. Recall however that the model also made specific predictions about the proportion of subjects that should make agent and patient interpretations for the different type of NPs in the experiment (see Figures 7 and 8). In this regard, it is important to note some properties about the activation of the interpretation nodes. The first is that the model predicts that subjects primarily make patient interpretations for good-agent and good-patient NPs with past participles (*arrested cop*, *arrested crook*), whereas they often make agent interpretations for good-

agent NPs with present participles (*arresting cop*). Interestingly, the model predicts that subjects will make a number of patient interpretations for NPs with present participles and typical patients (*arresting crook*).[8]

This last prediction was tested by examining the correlation between patient rating and response latency for the good-patient NPs with present participles (*arresting crook*) that were given patient interpretations by the model, and by examining whether the same correlation exists for the *arresting crook* NPs that were interpreted as agents by the model.[9] If the model correctly predicts which of these NPs subjects are giving patient interpretations in the on-line task, then there should be a significant negative correlation between patient rating and response latency. Importantly, there should not be a similar correlation for those items which the model predicts will be given agent interpretations. The rationale was that if subjects are giving the phrases patient interpretations, then they should be faster to do this as patient rating increases because it should be easier to overcome the strong bias of the present participle toward the agent role if the head noun is a highly typical patient.

The correlations supported the model's prediction. There was a significant correlation between patienthood rating and response latencies for the 37 items predicted by the model to be given patient interpretations in the on-line task $r = .37$. In contrast, the same correlation for the 16 items predicted by the model to be given agent interpretations was nonsignificant,

---

[8]Note that it is counter intuitive that the model makes this prediction because one would expect subjects to take longer to respond as the activation of the agent and patient interpretation nodes become more similar. However, this value represents the mean activation of the interpretation node. In fact, the activation of the patient interpretation node was either strongly activated (i.e., there was 35 NPs whereby the patient interpretation node was > .8) or hardly activated at all (there were 14 NPs whereby the patient interpretation node was < .2). Thus, there was only 4 NPs that fell between these extremes.

[9]The model's predictions for which of the items people were giving patient interpretations versus agent interpretations for the other NP types (i.e., *arrested cop*, *arrested crook*, and *arresting cop*) could not be tested because the model predicted, appropriately, that subjects made mostly one type of interpretation for these items.

$r = .19, p > .4$. In summary, the model not only predicts response latencies for the various conditions, it also predicts the NPs with present participles to which subjects are likely to give patient interpretations.

Discussion

Experiment 3 provided important evidence about how people actively combine morpho-syntactic information with knowledge of events to constrain thematic role assignment during NP interpretation. Specifically, the results show that the expectancy created by present participles leads to the facilitation in response latencies, and to an increase in the number of sensible judgements, when the head nouns were rated as typical agents for the participles in comparison to when they ware rated as typical patients. Moreover, the correlations indicated that subjects responded more quickly as head nouns became better agent role fillers for the present participles. The results for past participles demonstrated that they are biased toward focusing on patients in events, and that this bias combines with event knowledge to restrict the focus of the participles to event-specific patients. When the head nouns denoted typical patients, subjects responded more quickly and judged more of the items as sensible than when the same head nouns were paired with present participles.

Another important result was that grammatical marking and type of NP did not interact for unrelated trials. This indicates that the information being tapped in the on-line experiment is indeed detailed world knowledge of events, and that this information combines with morpho-syntactic information to constrain thematic role assignment during NP interpretation. If it were not event-specific knowledge that was driving thematic role assignment then agent and patient typicality of the head

nouns should have not influenced response latencies or percent sensible responses, and thus marking should not have interacted with NP type.

The results also indicated that the present and past participles differ in the strength to which they are biased toward the alternative roles. Although subjects consistently made patient interpretations for past participles, they did not always make agent interpretations for present participles. The ambiguous nature of the present participle was demonstrated by the negative correlation between patient rating and response latencies for *arresting crook* NPs. Thus subjects can easily generate patient interpretations for these NPs when the head nouns are very typical fillers for the patient role.

The pattern of correlations for good-patient NPs with present participles indicated that the model was predicting which of the NPs with present participles were given patient interpretations in the on-line task. One question that needs to be addressed is what is it about these NPs that invites subjects to make patient rather than agent interpretations. Visual inspection of the ratings for these items provides important information about why the model is interpreting these items as patients. All of these NPs are differentiated from the NPs given agent interpretations in the model because they are very polarized on ratings for the alternative roles of verbs. For example, the phrase *adopting baby* has a head noun that is very typical for the patient role, but is a terrible filler for the agent role. In contrast, the items that were given agent interpretations were those in which the filler was better for the agent role, but not as good for the patient role (e.g., *scratching gardener*). Another important property of the items given an agent interpretation were that some of these items were polarized on the different roles, but they all had patient head nouns that were not as

typical as the ones given patient interpretations. This suggests that it is not simply sufficient that the role/filler ratings be polarized on the alternative roles in order to be given a patient interpretation, but rather that they are polarized *and* rated as very typical patients. When this occurs, subjects can overcome the strong bias of the present participle to take the head noun as an agent, and more easily generate patient interpretations.

Recall that one alternative possibility for the large number of patient interpretations found for *arresting crook* NPs was the possibility that the participles' patient roles admitted a greater range of noun fillers, and thus subjects could generate these interpretations easily for present and past participles. One problem with this explanation is that one would have expected to find these biases in the unrelated trials as well. That is, if it were easier to fill the patient roles for the participles in general because they admitted a wider range of fillers, then subjects should have interpreted the unrelated trials faster when the participles were in their past than present forms. However, there was no evidence of this in the on-line data.

Note that a full discussion of the implication of these results for theories of conceptual combination is deferred until after Experiment 4.

3.8 Experiment 4

Although Experiment 3 provided insights into how people combine knowledge of events with grammatical morphemes, there were, unfortunately, some problems with the study that warranted a replication. One problem was that some NPs had to be removed from the analyses because their participles were not in their proper past participle form. A second problem was that a few participles and head nouns appeared twice

in each list. In Experiment 4, all past participles were in their correct form, and no subject saw any participle or head noun more than once.

Another difference with the current experiment was that the unrelated control trials were removed for two reasons. First, as discussed above, because the unrelated agent and patient NPs did not interact with participle marking in Experiment 3, it was clear that the information being tapped is event-specific information. Thus, it was unnecessary to have the unrelated control condition in Experiment 4 to show that it is event-specific information that is driving the combinatorial process. Second, the removal of the unrelated items permitted a stronger design because each subject received a greater number of items per condition.

Predicting On-line Modifier-Noun Interpretation in Experiment 4

Similar to Experiment 3, the goal of this section was to use the parameter settings obtained from the model that provided that best simulation of Experiment 2 to make predictions about how subjects would interpret these types of phrases on-line in Experiment 4. To derive these predictions, the 66 good-agent and 56 good-patient NPs used in Experiment 4 were used in a new simulation. Similar to the procedure used to make predictions about Experiment 3, each NP and its corresponding role/filler typicality ratings were entered into the model. The number of cycles of competition to reach criterion for each NP was recorded along with the activations of the alternative interpretation nodes.

The mean cycles of competition produced for each type of NP are illustrated in Figure 9. As illustrated in the figure, the model is making essentially identical predictions for this experiment as well. This was not unexpected, however, as there were 103 NPs that appeared in both

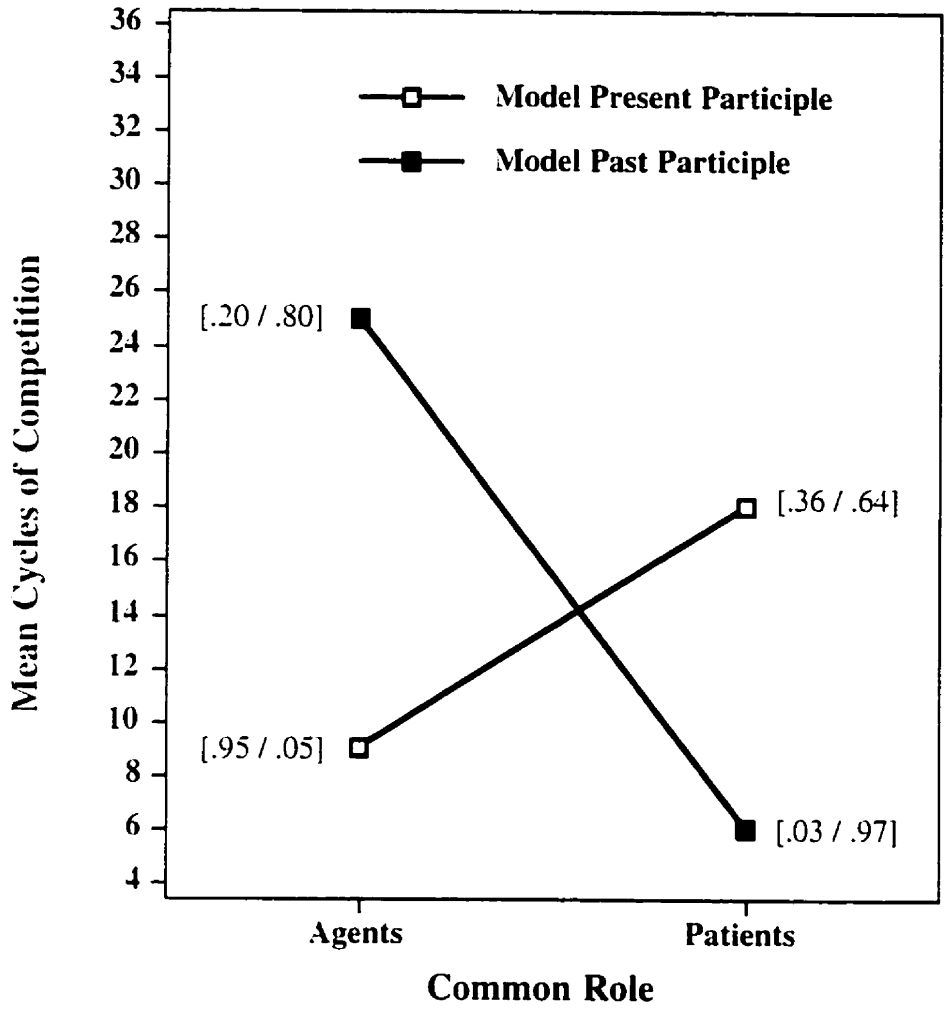

Figure 9. *The model's prediction for response latencies in Experiment 4. The values in brackets indicate the mean activation of the agent and patient interpretation nodes for each condition.*

Experiments 3 and 4. The only difference was that the model produced 1 cycle of competition less for good-agent NPs with past participles. The activation of the interpretation nodes were also similar.

Method

Subjects

Forty-four subjects participated for course credit, 11 per list. All subjects were native English-speaking psychology undergraduates from the University of Western Ontario, and all had normal or corrected-to-normal visual acuity.

Materials

Sixty-six phrases were created by pairing participles with a good-agent head noun, and another 56 noun phrases were formed by pairing participles with a good-patient head noun. These NPs included all 64 of the good-agent NPs and 56 good-patient NPs used in the off-line norming study. Similar to Experiment 3, the agenthood ratings of the 66 participle-agent pairs were significantly higher (M = 6.5, *range* = 5.7 - 6.9) than their patienthood ratings (M = 2.9, *range* = 1.3 - 6.9), $t_2(65) = 18.32$. In contrast, the patienthood ratings of the 56 participle-patient pairs were significantly higher (M = 6.0, range = 5.1 - 7.0) than their agenthood ratings, (M = 2.8, *range* = 1.0 - 6.6), $t_2(55) = 15.71$.

Four lists were constructed that included 2 with NPs with good-agent head nouns (*arresting cop*), and 2 that consisted of NPs with good-patient head nouns (*arrested crook*). Each of the good-agent NP lists included 33 NPs with present participles, 33 NPs with past participles, and 66 filler trials that consisted of nonsensical participle-noun NPs (e.g., *brewed ant*). Alternatively, each of the lists that contained good-patient NPs had 28 NPs with present participles, 28 NPs with past participles, and 56 filler trials

that consisted of nonsensical participle-noun NPs. Across the different lists, each participle appeared in both its present and past form, and modified a related head noun. No subject saw any participle or head noun more than once. Finally, 20 practice trials were created that consisted of 10 with present participles and 10 with past participles. Half of these items were nonsensical and half were sensible.

Procedure

All aspects of the procedure were identical to Experiment 3.

Design.

Response latencies and percent sensible responses were analyzed by three-way analyses of variance. The factors of interest were grammatical marking (present vs. past), which was within subjects and items, and NP type (good-agent vs. good-patient) which was between subjects and items. List and rotation group was used as a between item variable to stabilize variance due to rotating subjects and items across different lists. Similar to Experiment 3, correlational analyses were also performed. Note that a trial was excluded from the latency analyses if the response indicated it was nonsensical.

Results

Analyses of Variance

Decision latencies. Response latencies greater than 3 standard deviations from the mean were replaced by that value (less than 1% of the scores). Response latencies are presented in Table 8.

Grammatical marking and role interacted, $F1(1, 40) = 13.84$, $F2(1, 118) = 12.83$. Planned comparisons revealed that good-patient NPs were responded to more quickly when they were in their past than present form,

Phrase Type	*M*	*SE*
Good-Agent		
(-ing)	1646	99
(-ed)	1813	132
	167**	
Good-Patient		
(-ing)	1710	81
(-ed)	1552	67
	158**	

**Significant by subjects and items

Table 8

Mean response latencies (ms) for good-agent and good-patient NPs in Experiment 4.

Phrase Type	*M*	*SE*
Good Agent		
(-ing)	89	2.2
(-ed)	65	4.9
	**24	
Good Patient		
(-ing)	80	3.5
(-ed)	92	1.1
	**12	

**Significant by subjects and items

Table 9

Percentage of good-agent and good-patient NPs judged sensible in Experiment 4.

$F1(1, 40) = 6.51$, $F2(1, 118) = 7.82$. Alternatively, good-agent NPs were responded to more quickly in their present than past form, $F1(1, 40) = 7.27$, $F2(1, 118) = 5.07$. The main effects of role and marking were not significant, all F's < 1.

Finally, it should be noted that overall subjects responded relatively slowly when making "nonsense" judgements (M = 1990 ms, SE = 18 ms).

Percent sensible responses. Percent sensible responses are presented in Table 9. Grammatical marking and role interacted, $F1(1, 40) = 48.29$, $F2(1, 118) = 133.51$. Planned comparisons revealed that good-patient NPs were responded to as sensible more often when they were in their past than present form, $F1(1, 40) = 10.63$, $F2(1, 118) = 27.24$. Alternatively, good-agent NPs were responded to as sensible more often in their present than past form, $F1(1, 40) = 42.52$, $F2(1, 118) = 128.43$.

There was a main effect of role, in that patient items were responded to as sensible more often (M = 86%, SE = 2.0%) than agent items (M = 77%, SE = 3.2%), $F1(1, 40) = 6.08$, $F2(1, 118) = 19.06$. There was also a main effect of marking; NPs with present participles were judged sensible more often (M = 85%, SE = 2.1%) than NPs with past participles (M = 79%, SE = 3.2%), $F1(1, 40) = 5.14$, $F2(1, 118) = 14.21$.

Correlation Analyses

The results of the correlation analyses for response latencies are shown in Table 10, whereas Table 11 contains the results for percentage of sensible responses.

Response Latencies. The pattern of correlations were similar to Experiment 3 in a number of ways. First, there was a significant negative correlation between agent rating and response latencies for good-agent NPs with present participles (*arresting cop*), demonstrating that subjects'

Phrase Type	*Agent Rating*	*Patient Rating*
Good-Agent		
(-ing)	-.36**	-.10
(-ed)	-.08	-.25**
Good-Patient		
(-ing)	-.12	-.31**
(-ed)	.09	-.36**

**indicates a significant correlation

Table 10

Correlations between patient and agent ratings and response times for all NPs in Experiment 4.

Phrase Type	*Agent Rating*	*Patient Rating*
Good-Agent		
(-ing)	.25**	-.14.
(-ed)	-.30**	.44**
Good-Patient		
(-ing)	.19	.33**
(-ed)	-.14	.17

**indicates a significant correlation

Table 11

Correlations between patient and agent ratings and percentage of sensible responses for all NPs in Experiment 4.

response latencies decreased as the agenthood ratings of the head nouns increased. Second, there was a significant negative correlation between patienthood rating and response latencies for good-patient NPs with past participles (*arrested crook*). Thus, similar to Experiment 3, subjects were sensitive to thematic fit of head nouns of NPs in which the morphemes and head nouns were consistent.

When the fillers and head nouns were inconsistent (*arrested cop*, *arresting crook*), subjects were again sensitive to the typicality of the head nouns. For *arresting crook* NPs, there was a negative correlation between patienthood rating and response latencies. This result is explored further below. One difference between Experiments 3 and 4 was a significant negative correlation between patienthood rating and response latencies for good-agent NPs with past participles (*arrested cop*), indicating that subjects responded more quickly as the patient rating of the head nouns increased. Although a similar correlation was found in the percentage of sensible responses in Experiment 3, the same correlation in the response latency data was not significant. To investigate whether this difference was simply due to different items in the two experiments, or whether subjects were responding to the same NPs differently in Experiment 4 than in Experiment 3, the correlation was performed again on only the 52 good-agent NPs that appeared in both experiments. This correlation was now non-significant (-.15, $p > .29$), and similar to Experiment 3 (-.11). No other correlations for response latencies were significant.

Percent sensible responses. Similar to Experiment 3, there was a negative correlation between percent sensible responses and agenthood rating for *arrested cop* NPs, and a positive correlation between patienthood rating and percent sensible interpretations for these NPs. Unlike

Experiment 3, the positive correlation between percent sensible responses and agenthood rating for *arresting cop* NPs was significant.

Another difference between the two experiments was the pattern of correlations for the good-patient NPs with present participles (*arresting crook*). Unlike Experiment 3, there was a significant positive correlation between patienthood rating and percentage of sensible responses. Furthermore, the correlation between agenthood rating and percentage of sensible responses for these items was now nonsignificant. Note that these differences are partly accounted for by the different items in the experiments. For the overlapping items, the correlation between patient rating and percentage of sensible responses was now nonsignificant (.23, $p > .12$), although there was a trend for subjects to indicate that these NPs were more sensible as the head nouns increased in patienthood. Alternatively, the agent ratings of these NPs also approached significance (.23, $p > .12$), indicating that for exactly the same items used in Experiment 3 subjects were finding them less plausible as they increased in agent rating. The differences observed between the correlation analyses for Experiments 2 and 3 are discussed in more detail below in the Discussion.

Model Predictions

Figure 10 contains the model's predictions along with the corresponding human response latencies. The model again captures all major aspects of the response latency data. Specifically, when the modifiers were past participles the model made the correct prediction that subjects should take longer to respond when the head nouns were typical agents (*arrested cop*) than when they were typical patients (*arrested crook*). Alternatively, when present participles were used, the model predicted the relative ease that subjects had when interpreting NPs with typical agent

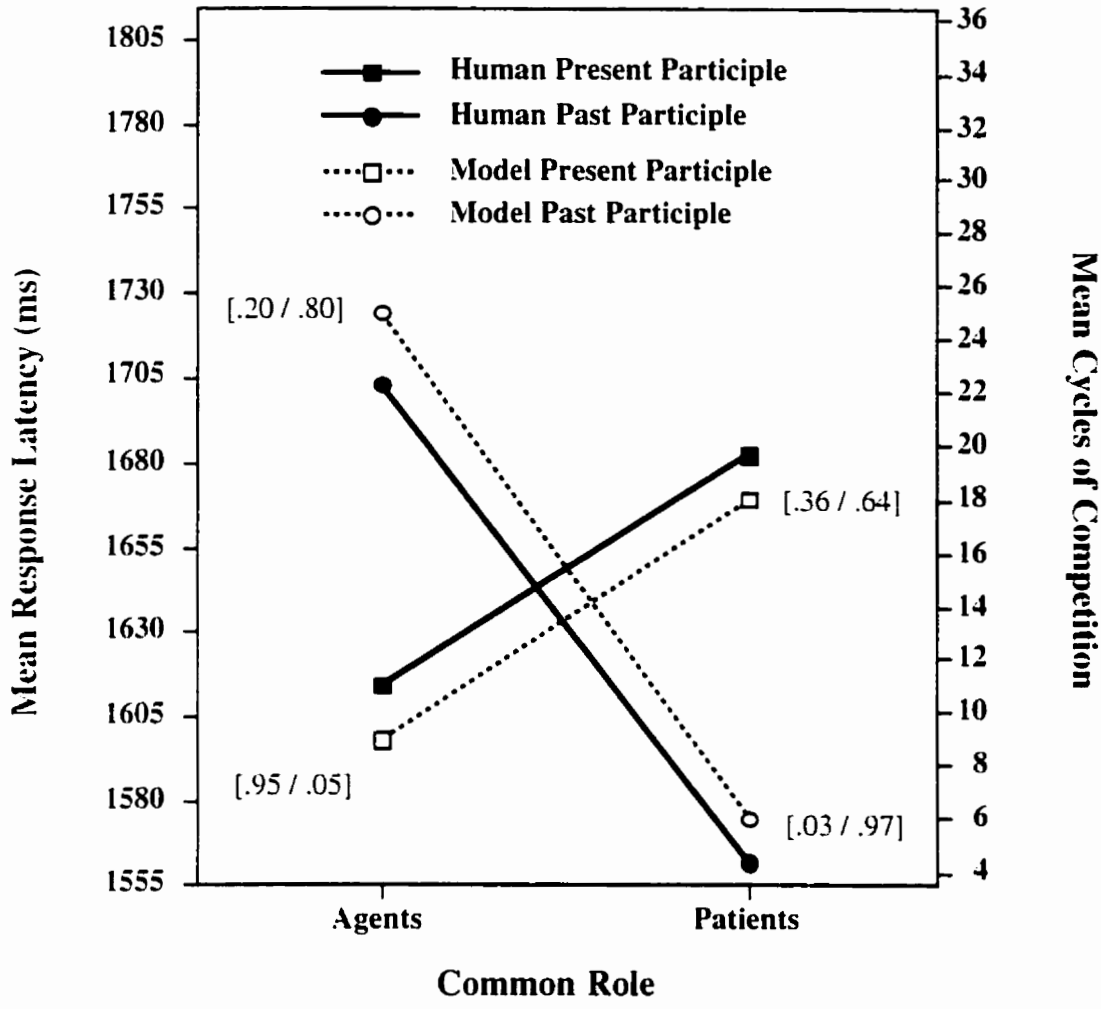

Figure 10. *Experiment 4 response latencies presented along with model's prediction for the same conditions. The values in brackets indicate the mean activation of the agent and patient interpretation nodes for each condition.*

(*arresting cop*) versus typical patient (*arresting crook*) head nouns.

The correlations examining the model's predictions concerning which of the good-patient NPs with present participles were given patient interpretations on-line were similar to those reported in Experiment 3. Specifically, there was a significant correlation between patienthood rating and response latencies for the 38 NPs given patient interpretations by the model (.35), whereas the same correlation for the 18 NPs interpreted as agents by the model was non significant (.21, $p > .4$). Thus, the model appears to select which of these NPs subjects interpret the head nouns as patients rather than as agents.

Discussion

The main goal of Experiment 4 was to replicate and extend the findings of Experiment 3. Experiment 4 showed that subjects responded more quickly, and found more NPs sensible, when the head nouns were good fillers for the expected role of the participles. Moreover, subjects responded more slowly, and interpreted fewer NPs as sensible, when the head nouns were poor fillers for the expected role of the participles. As illustrated in Figure 10, the model's prediction for the response latencies in the different conditions were closely captured in the on-line data.

Despite the many similarities between the results of Experiments 3 and 4, there were some differences. The first difference was the percentage of *arrested cop* NPs that were judged as sensible. In Experiment 4, only 65% of these NPs were judged sensible, whereas in Experiment 3, 85% of these types of NPs were judged sensible (note that this held also for the items that overlapped). These results are consistent with the stronger negative correlation between agent rating and percent sensible responses in Experiment 4. However, this influence appears to be limited to percent

sensible responses because the same correlation in the response latency data was similar to Experiment 3 when only the overlapping items were compared.

The second difference was more subtle, but nonetheless present in the pattern of results found in the percent sensible response correlational analysis for good-patient NPs with present participles (*arresting crook*). Specifically, subjects in Experiment 4 were finding more of these NPs sensible as they increased in patient rating, and fewer of them sensible when they were better agents. Although removing the NPs that did not overlap in the two experiments accounted for some of the difference between the two experiments, it did not eliminate them. Note however, that the correlation analysis in the response latency data was similar to Experiment 3; the only exception was that correlations were stronger. That is, agent and patient ratings of the head nouns had a stronger influence on response latencies for the same items in the two experiments. Taken together, these differences suggest that in Experiment 3 subjects were more sensitive to the typicality of the head nouns for the alternative roles of the participles, and that these differences are most apparent in the percentage of sensible responses made by subjects for good-agent NPs. The similar pattern of response latencies across the two experiments is also evidenced by how well the competition model predicted the obtained data.

One plausible reason for the difference in the percent sensible responses could be that subjects in Experiment 4 received no unrelated trials. Thus, for half of the trials subjects were interpreting related NPs in which the head noun was always a good filler for at least one of the participles thematic roles, and for the other half of the trials the head nouns were nonsensical for both roles (e.g., *brewing ant*). It is possible that

removing plausible, but not typical fillers, contributed to the differences between the two studies. For example, the contrast between what head nouns are good fillers (therefore, sensible) and those that are not (nonsensical) is sharper for subjects in Experiment 4 because they judged fillers that were either terrific fillers for one of the participles roles, or very implausible (*brewing ant*). This sharp contrast may be causing subjects to be more sensitive to both the patient and agent typicality of the head nouns. It could be this sensitivity that is leading to the stronger correlations for patient rating and percent sensible responses for *arresting crook* NPs, and for fewer of *arrested cop* NPs to be judged sensible.

3.9 Implications for Theories of Conceptual Combination

The research reported herein has a number of implications for theories of conceptual combination (e.g., Gagné; 2000; Gagné & Shoben, 1997; Murphy, 1988; 1990; Springer & Murphy, 1992; Wisniewski, 1996). In this section, the theoretical approaches to conceptual combination are first discussed, and then the research examining how people interpret noun phrases with past participle modifiers is reviewed and evaluated with respect to the current research.

There are three main models of conceptual combination: the schema modification model (Murphy, 1988; 1990); the dual-process model (Wisniewski, 1996; Wisniewski & Love, 1998); and the competition among relations in nominals (CARIN) model (Gagné, 2000; Gagné & Shoben, 1997). Each is discussed in turn.

Schema Based Theories. The two schema-based theories of conceptual combination (schema modification and dual-process) both follow from the assumption that concepts are organized into structured lists of slots and

fillers (Minsky, 1975; Rumelhart, 1980). As discussed in Chapter 1, according to schema based approaches, slots represent the general dimensions of a concept along with their default values. For example, an entity such as a dog may have a slot that specifies its habitat, which may have a range of values such as home, farm, or streets. In this case, the default value for the habitat slot would be home because it is the most typical habitat of dogs.

According to the *schema modification model* (Murphy, 1988; 1990), comprehension of modifier-noun phrases (e.g., adjective-noun and noun-noun combinations) involves two-stages. In the first, the modifier selects and fills a slot in the head noun. Selecting a relevant slot to modify can occur in a number of ways. The simplest occurs when the representation of the modifier previously exists in the representation of the head noun, and so selecting the appropriate slot is relatively straightforward - slot filling involves matching the modifier to the appropriate slot in the head noun (e.g., the color slot in apple for the combination *red apple*). The second method involves narrowing the range of appropriate slots by using a constraining context. Thus, the appropriate slot for a modifier to fill is determined by the context in which the combination is embedded. The third way a slot can be chosen is by using general knowledge of the world. For example, to understand *apartment dog,* people use their world knowledge to realize that the "habitat" slot is the most appropriate because dogs are more likely to inhabit an apartment then they are to bite apartments or look like apartments (Cohen & Murphy, 1984; Murphy, 1988). It should be noted that not all variants of the schema modification approach allow for world knowledge to have an immediate role in constraining slot selection. For example, according to the selection modification model (Smith, Osherson,

Rips, & Keane, 1988), the selection of the appropriate slot to modify in the head noun is driven initially by a process that is akin to spreading activation between constituent concepts. Thus in this initial stage of processing, no information outside of the constituent concepts is used.

Following the process of selecting and filling the appropriate slot, general knowledge outside of the two concepts may be used further to elaborate the combination so that it is more plausible (Murphy, 1988; 1990). Take the *apartment dog* example again. After selecting the "habitat" slot people may use their world knowledge about apartment dogs to elaborate the combination by inferring that apartment dogs are more likely to be well behaved or are small in comparison to other dogs. The important point is that people will use their knowledge of apartments, dogs, and how the two interact to refine their representation of the combination (Murphy, 1988).

The evidence considered as support for this model comes from a number of sources showing that combinations are comprehended faster if the modifier can easily select a slot in the head noun (Murphy, 1988; 1990; Smith & Osherson, 1984), and from studies illustrating that combined concepts involve features that are not part of the constituent concepts (e.g., Murphy, 1988; 1990; Springer & Murphy, 1992).

The *dual-process theory* retains many of the representational assumptions of the schema modification model, but emphasizes that concepts often change considerably during conceptual combination in ways that cannot be accounted for by slot-filling or relational linking (Wisniewski, 1996; Wisniewski & Love, 1998). Specifically, according to this perspective, people often construe a noun to be a representation of something associated with the noun rather than the noun itself. That is, people will construe a representation of a noun by utilizing conceptual

information related to its referent such as a property or a thematic associate (Wisniewski, 1997). For example, the combination *moose pencil* may be interpreted as "a pencil with an eraser that looked like a moose". In this case, the referent for the noun is not an actual moose (i.e., the thousand pound creature that stinks), but rather a construed representation of a real moose.

Wisniewski proposes comparison and construction processes to account for the construal of nouns in the manner described above. These processes are assumed to operate simultaneously in a highly interactive manner. During the comparison process, people align or compare the structure of the constituent concepts, and it is by this process that the commonalties and differences between the concepts become apparent (Wisniewski, 1997). Finding differences between the concepts is important for determining how the head noun is modified. Similar to the schema modification model, the process of selecting differences between concepts is constrained by plausibility derived from world knowledge. Wisniewski suggests that the construction process accounts for people's ability to construct a novel property in a combination that is not part of the constituent concepts. The similarities and differences highlighted by the comparison process play a large role in constraining how a property of the modifier should be integrated into the head noun. The construction process works with this knowledge to construct a novel representation of the head noun, while at the same time is heavily constrained by world knowledge to maintain the "representational integrity" of the constituent concepts.

Another aspect of this approach that is relevant to the current discussion is the proposal that scenarios (i.e., situations denoting states, events, and processes) are part of a noun's schema representation and are

used to constrain relational interpretations. For example, the appropriate relation for the combination *truck soap* (i.e., soap for cleaning trucks) would be constructed easily because a plausible "cleaning" scenario (as evaluated by world knowledge) can be applied in which the truck fills the patient role and soap fills the instrument role. This occurs because the concept "soap" makes available information about a cleaning scenario that involves multiple roles (i.e., agent, patient, and instrument) and potential role fillers. Alternatively, combinations with similar constituents such as *magazine newspaper* are proposed to be more difficult to interpret by a relational interpretation because both constituents are appropriate fillers for the same role (e.g., the patient role of read), and neither are appropriate for an alternative role (Wisniewski, 1997).

Evidence for the comparison and construction processes has been provided by studies showing people sometimes interpret novel combinations in ways other than slot-filling or applying a general relation, such as property and hybrid interpretations (e.g., Wisniewski, 1996; Wisniewski & Love, 1998; Wisniewski & Gentner; 1991; but see Gagné (2000) for contradictory research). Unfortunately, direct evidence for scenario construction in the manner described by Wisniewski has not been shown. However, this approach does make the suggestion that event knowledge associated with nouns play a crucial role in specifying which relations should be drawn between constituents in events. In this regard, the present Experiments 2, 3, and 4 are the first to provide direct support for these suggestions.

The Competition Among Relations in Nominals (CARIN) Model. According to the CARIN model (Gagné, 2000; Gagné & Shoben, 1997), comprehending combinations begins by applying a general thematic

relation between the two constituent concepts (e.g., location relation for the combination *apartment dog*), followed by the use of world knowledge to narrow and evaluate the plausibility of the relation. Importantly, the relation that is applied is the one with the greatest strength, which is indexed by the relative frequency of a relation for a noun when it is used as a modifier. If multiple relations are frequent for the noun, then the selection of the relation with the highest strength will be more difficult than when there are only a few because the relations compete for selection. Thus, how a concept is used linguistically (i.e., distribution of relations for a noun when used as a modifier) is known and utilized by the comprehender to guide interpretations.

This model differs from the schema models in a number of ways. First, the CARIN model proposes that interpreting modifier-noun combinations proceeds by first selecting a general relation that constrains how world knowledge will be subsequently employed to elaborate the combination. Second, it is the relations associated with the modifier rather than the head noun that are important for guiding interpretation. Gagné and colleagues emphasize that modifiers typically appear before head nouns in English, and are used to denote an instance that belongs to the category of the head noun. Third, CARIN borrows notions from linguistic accounts of nominal compounds (e.g., Kay & Zimmer, 1976, Levi, 1978) which emphasize general thematic relations (e.g., about, made of, causes, etc.) rather than more specific dimensions of schemas (e.g., habitat slot or color slot). Thus according to the CARIN model, *apartment dog*, is interpreted by applying a general location relation rather than modifying the more specific "habitat" dimension in "dog."

Support for this theory has been provided by Gagné & Shoben (1997) and Gagné (2000). Gagné and Shoben demonstrated that phrasal judgement times are faster when the relation used to interpret the constituent concepts is highly frequent of the modifier, but whether or not it is frequent of the head noun does not matter. Furthermore, regression analysis show that combinations take longer to be judged sensible when they include modifiers with numerous frequent relations.

Adjectival Participles and Conceptual Combination

Adjectival participles play an important role in the conceptual combination literature. Specifically, they have been used in studies illustrating that combined concepts involve features that are not part of the constituent concepts, and more importantly, that these features (called phrase features) become available before features of the head noun (e.g., Gagné & Murphy, 1996; Potter & Faulconer, 1979; Springer & Murphy, 1992). Springer and Murphy (1992) collected sentence verification times for properties of phrases containing past participles such as *peeled apple* and *boiled celery*. Subjects received four versions of verification sentences, including those in which the properties were true of the head noun (e.g., Peeled apples are round / Boiled celery is green), true of the phrase (Boiled celery is soft / Peeled apples are white), false of the noun (Boiled celery is blue / Peeled apples are square), and false of the phrase (Boiled celery is crisp / Peeled apples are red). Subjects were faster to verify features that were true of the phrase versus those that were true of the head noun, and verified features false of the phrase and head noun equally quickly.

These results contradict one of the main assumptions of the selection modification approach to adjective-noun interpretation (Smith, Osherson, Rips, & Keane, 1988). Recall that according to this approach, people

interpret adjective noun phrases through a two-stage serial process. During the first stage, the model holds that fast dimension selection occurs via spreading activation from the constituent concepts. If the adjective easily fills a dimension in the head noun schema (e.g., celery is green), then people should be quick to verify the feature. The second stage involves a slower world knowledge-dependent process in which other features of the combined concept are inferred if necessary to form the novel concept. Thus, according to this perspective, knowledge outside of the constituent concepts is not used during the initial stage of processing. The fact that phrase features were verified more quickly than noun features indicates clearly that it cannot be the case that the constituent concepts are accessed first and then combined. Springer and Murphy (1992) account for these results by suggesting that when people interpret phrases such as *boiled celery* they do not activate all conceptual information associated with celery, but rather they only consider what happens to celery when it is boiled. This proposal also explains why features false of the phrase and head nouns were verified equally quickly. For both of these sentences, the features are false of the phrase, and because people focus on the phrase features, they do not respond faster for a feature true of the head noun but false for the phrase.

Springer and Murphy's results are also consistent with the schema modification model (Murphy, 1988; 1990) and the dual process model (Wisniewski, 1996) because both models hold that world knowledge is employed initially to constrain relational interpretations. Applying the schema modification approach to the combination *peeled apple*, one would expect that the modifier selects a slot or dimension in the noun apple to modify, and that this initial process involves world knowledge that is only relevant to what happens when apples are peeled. From this world

knowledge, a particular dimension in the head noun (perhaps color) would be selected and modified so that it captures the phrase feature that peeled apples are white. One problem with this approach is how world knowledge is constrained by the linguistic properties of the modifier is not clear. For this reason, it is difficult to know how world knowledge about events such as peeling apples is used to draw the intended relation.

The assumptions of the dual-process theory are very similar to the schema modification approach. The improvement with this model is the proposal that scenarios associated with the constituent concepts could provide important information about the intended relation between them (Wisniewski, 1997). For example, when people read *peeled apples*, they can use their knowledge about peeling scenarios to infer that the head noun is most likely to fill the patient role, and presumably evaluating the fit of the filler for the thematic role could constrain activation to features that are only consistent with peeled apples - such as the fact that peeled apples are white. This proposal is discussed further below.

The model that has made the most explicit assumptions regarding the processing of phrases such as peeled apples is the CARIN model. Gagné and Murphy (1996) suggest that the results could be accounted for by first identifying a linguistic relation between the modifier and the head noun, and then using this relation to infer the relevant features for the combination. For example they state:

> "readers first understand peeled apple as an apple that has had its peel removed (i.e., inferring the intended relation between peeled and apples). However, they do not access all of their conceptual knowledge about apples at this point. Next to determine which features should be included with the new concept, they consider the implications of the particular relation that is selected. That is, they consider how "an apple that is peeled" differs from apples that are not peeled. For example, peeled apples no longer have skins, and, therefore they are not red (or green) like other

apples. The features that distinguish peeled apples from apples are the first to become available in the new concept. Identifying these features requires one to refer to world knowledge or examples that one has encountered of the concept, or both. Also, this attempt to identify distinctive features of peeled apples is consistent with our observation that nonredundant information is verified more quickly." (p. 99 - 100).

Thus all three of the main models can account for the results of NPs with past participles such as peeled apples. All models would hold that this is accomplished via relation linking between the head noun and participle. In this regard, the schema modification model is probably the most underspecified because it is not clear how people draw the intended relation between the constituent concepts. That is, the relation is associated with the modifier rather than the head noun. Both the Dual Process and CARIN models make more explicit assumptions about how the relation is drawn. These models are now discussed in relation to the current data.

Extending Theories of Conceptual Combination

The question addressed in this section is how the data presented in Chapter 3 are accounted for by models of noun phrase interpretation, and to compare these models with the schema based approach proposed in this chapter. Similar to the results found in the literature examining how people interpret phrases such as *peeled apples*, the results for the noun phrases in the current chapter also cannot be easily accounted for by the selection modification model (Smith et al., 1988). The main problem for this approach is accounting for the interaction between noun phrase type and grammatical marking. Specifically, an initial stage of processing involving spreading activation from participles and head nouns should lead to the activation of information about typical agents and patients in events. Therefore, the model fails to predict the grammatical marking by NP type interaction found in Experiments 3 and 4. Thus, the same problems

associated with spreading activation theory that were discussed in Chapter 2 are also applicable here. Recall that the most difficult problem for such approaches is to explain how information activated via spreading activation is constrained by grammatical morphemes.

The results could be accommodated by the two schema based approaches if they were augmented to deal with the focusing properties of grammatical morphemes and the argument structures of participles. In their current state, they are preferable to the selection modification model because they allow for world knowledge about the roles participants play in specific events to constrain the initial drawing of appropriate relations (see Wisniewski, 1997). The problem for the current schema approaches is explaining how world knowledge combines with grammatical morphemes and the argument structures associated with the modifiers to constrain relation interpretation. However, these approaches could easily be augmented to allow for an interaction between thematic relations and grammatical morphemes. However, by doing this the models would place additional emphasis on the role of thematic relations associated with modifiers, thus making them more similar to the CARIN model.

The schema based model proposed in this chapter is similar to the CARIN model in many ways. For example, one of the main assumptions of the CARIN model is that thematic relations compete for modification of the head noun. Recall that in this model the selection of a relation to modify a head noun results from a competition process in which the relation with the greatest strength is selected. When there are multiple strong relations, there is more competition which leads to longer interpretation times. Alternatively, less competition occurs when one relation is far stronger than the other(s), which leads to shorter response latencies. The

assumption of competing thematic relations associated with modifiers is also an explicit assumption of the model proposed herein.

The two models also share other important representational and processing assumptions. Recall that thematic role information in the current model is proposed to be represented in a form similar to connectionist and instance-based perspectives of schema processing (e.g., Golden & Rumelhart, 1993; Hintzman, 1986; Rumelhart et al., 1986; Sharkey & Sharkey, 1992). In many ways, this approach to thematic role information is similar to the CARIN model which assumes that people represent and use distribution information about the frequencies of relations for modifiers during NP interpretation. The similarities between this form of representation and connectionist perspectives of knowledge representation is well recognized by Gagné and Shoben (1997) who stated that:

> "....our assumptions about the use of distributional information are congruent with recent connectionist accounts of language processing (e.g., Elman, 1990; Seidenberg, 1992) that are sensitive to probabilistic information at many levels of the linguistic representation as a natural consequence of their learning algorithms. In addition, our assumptions about the use of thematic relations is not unlike the information about use that is commonly postulated to be stored with verbs (MacDonald, 1994; MacDonald et al., 1994; Trueswell et al., 1993)" (p. 83).

Thus the two models make similar proposals regarding the organization of thematic relations in memory. The main difference is that unlike the CARIN model and previous research on thematic role processing, McRae et al.'s (1997) view of thematic roles as verb-specific concepts includes the notion that thematic relations often involve specific conceptual knowledge. For example, McRae et al. argued that the thematic role of agency is not always represented in terms of general knowledge such as <animacy>, but

rather for verbs with well-defined role concepts, such as *arrest*, this information can include knowledge that the agent of *arrest* is <authoritative> and <powerful>. The detailed nature of this information is more specific than that proposed by the CARIN model, but is consistent with the two schema based models' assertion that more specific world knowledge is used to constrain dimension (i.e., slot/role) selection and filling during conceptual combination. Furthermore, this detailed knowledge is necessary to account for the results of Experiments 2, 3, and 4.

The following example illustrates how the CARIN model fairs as an account of the present results. According to the CARIN model, interpreting a phrase such as *arrested crook* begins by applying a general linguistic relation associated with the participle modifier. In this first stage people may first recognize that the most frequent relation associated with the modifier is an event relation, rather than other general linguistic relations such as "made of" or "has". Once people have utilized this general information to establish that an event interpretation is appropriate they then can use more specific information regarding the type of event relation. This relation should be a patient relation, because past participles are used most frequently to modify the patient role of head nouns. Because this model holds that thematic relations are general in nature, the patient relation would be limited to features that are consistent across most events, such as <undergo a change of state> (see Dowty, 1991; and Schlesinger, 1995 for lists of general patient features). Thus the head noun would initially be distinguished from other members of the head noun category (crooks), in that this "crook" undergoes a change of state. At this point the particular relations should not include any event-specific knowledge about people that are commonly arrested. Following this initial process, world knowledge

about how typical the filler is for the patient role of the specific event denoted by the participle is used to compute the fit of the head noun for this relation, and perhaps embellish the representation with detailed knowledge such as the crook is probably located in a jail cell. If the head noun is a good filler, then no additional knowledge is needed to develop a sensible interpretation and people can respond quickly. If it is a poor filler, then people must infer additional knowledge to make the phrases sensible, and thus be slower to respond. In the same manner, the CARIN model can account for the relative judgement times and percent sense responses for the unrelated items by holding that following the selection of an event relation, they can then use more specific information about the type of agent and patient relations for present and past participles, and because unrelated items are equally good fillers for the opposite roles, they take about the same amount of time to be interpreted. However, because they are not typical for either role it takes longer for the subject to respond than for related items because they must construct an interpretation based less familiar scenarios.

The discussion above indicates that the CARIN model can account for the current data without making the assumption that event-specific knowledge initially drove the application of the applicable relation but rather had an influence during a second stage of processing involving detailed world knowledge. It is important to note however, that an explanation based on general relations alone is not sufficient, as the data clearly show that response latencies are ultimately determined by how good the head nouns are as participants in the specific events denoted by the participles. Thus, it is possible that a general relation was drawn first but there was no evidence suggesting that this information was directly responsible for the observed data.

One aspect of the data that is a little more problematic for the CARIN model is how it would account for the pattern of data found for good-patient NPs with present participles (*arresting crook*). The model and the on-line and off-line experiments suggested that people made patient interpretations for some of these NPs, and they were faster to respond when they did. In contrast, when the head noun was a good agent, people almost always made agent interpretations (*arresting cop*). It is not clear how this could occur if it was thematic relations associated with the modifiers that dictate which relations are selected. Clearly in these examples it is information associated with the head noun, and not just the modifier, that is dictating which thematic relation is ultimately drawn between the head noun and modifier. Further, the pattern of correlations indicate that event-specific typicality of the participants is contributing to the ease with which people are drawing these relations. Thus, perhaps the CARIN model would have to be augmented to allow for event-specific information associated with the head nouns to drive selection of the appropriate relations, rather than holding that a general relation is first selected and then used to drive how world knowledge will be subsequently employed to elaborate the combination. In other words, in its current form the CARIN model emphasizes that information associated with a modifier determines the relation that is drawn. The research reported above indicates that this approach would have to be extended so that it can account for the role of the head noun in driving relation selection. One possibility is that people are more likely to utilize information associated with the head noun to determine relation selection when the relations associated with the modifier become more ambiguous (i.e., the frequencies of the relations become more even). Finally, it is important to note that the CARIN model

has to date been applied only to general thematic relations such as "made of" and not to thematic roles such as agent and patient. Thus, it is more appropriate to view the current research as an extension rather than as an replacement for the CARIN model.

In conclusion, the research reported in this chapter extends theories of modifier-noun phrase processing by making it explicit how people utilize knowledge of events to constrain noun phrase interpretation. The results also extend theories by showing how morpho-syntactic information associated with modifiers is used to constrain activation of world knowledge about events. Moreover, a major dimension along which theories of conceptual combination differ is the extent to which they emphasize the modifier (i.e., CARIN) versus the head noun (Schema Modification and Dual Process models). Although it may be possible to identify situations or types of noun phrases for which the modifier versus head noun are differentially salient for deriving interpretations, the data and modeling demonstrate that information associated with both the modifier and head noun is important. Finally, in line with the overall purpose of this chapter, these results suggest that people actively integrate the focusing properties of grammatical morphemes with world knowledge of events to constrain on-line thematic role assignment.

Chapter 4: General Discussion

The goal of this research was to investigate how people combine morpho-syntactic information with event-specific knowledge when interpreting phrases in isolation. This research made four main contributions. First, Experiment 1 demonstrated that locations and instruments of specific events are more activated when verbs reference situations as ongoing (imperfective aspect) versus completed (perfect aspect). Thus, given the brief SOA, verb aspect quickly modulates the activation of world knowledge of events when verbs are read in isolation. These results are particularly problematic for theories of semantic priming that are based on the notion of spreading activation because they do not include mechanisms that can account for the modulating influence of verb aspect. Although additional post hoc assumptions could be incorporated, such as adding the assumption that the initial spread of activation is constrained by morpho-syntactic information, the new spreading activation approach would no longer be in the same spirit upon which the notion of spreading activation was based. Finally, it is worth noting that Experiment 1 represents the first time in the semantic priming literature that verbs with different aspect morphemes have been used as primes in a short SOA priming task.

A second major contribution is that Experiments 2, 3, and 4 show how people actively integrate world knowledge of agents and patients in events with the aspectual properties of adjectival participles to constrain off-line and on-line thematic role assignment during NP interpretation. To my knowledge, these results are the first to demonstrate this combinatorial process during language comprehension. Moreover, the contribution of this research was enhanced by an implemented competition model that made

this combinatorial process explicit, and was used to generate specific predictions about exactly which NPs were given agent and patient role assignments on-line. According to this model, the process of interpreting noun phrases with participle modifiers involves competition between thematic role assignments for head nouns. How difficult it is to resolve competition depends on how biasing the inflections are toward the alternative thematic roles, and how good the head nouns are as fillers for the alternative roles of the participles. When multiple constraints support the same interpretation, people more easily generate a sensible interpretation, whereas when the support is more balanced for the two competing alternatives, it takes longer for people to generate sensible interpretations. The results of the on-line and off-line experiments provide strong support for the model, and extend current models of conceptual combination by making explicit assumptions about how people use their world knowledge of situations and morpho-syntactic information to constrain combinatorial processing. Although the model is a valuable contribution to the literature, there are nonetheless some limitations of the model that should be recognized. One limitation is that the model captures only how quickly and how often people interpret the phrases as sensible. That is, the model does not make nonsense interpretations like subjects sometimes do in the on-line task. A second limitation is that the model examines only the influence of thematic fit and participle bias. There are, however, other sources of information that may play a role during interpretation of these types of phrases, such as the frequency with which the participles appear as verbs versus adjectives in language.

The results reported in Chapters 2 and 3 also contribute to the psycholinguistic research examining how verb aspect and other

grammatical morphemes function to constrain the construction of situation models from events mentioned in text (Carreiras, Carriedo, Alonso, & Fernandez, 1997; Morrow, 1985a; 1990; Morrow, Greenspan, & Bower, 1987). The current results show that verb aspect does not only modulate the activation of information explicitly mentioned in the text, but also that it plays a role in foregrounding and backgrounding world knowledge about the common properties of events. The profiling of world knowledge specific to the temporal and causal components of the events makes this information potentially constraining for guiding expectations about the events. Thus, these results show that grammatical morphemes play a crucial role in distinguishing between relevant and irrelevant event knowledge associated with events mentioned in the text. Clearly, determining the level of detail of this information is an important step toward understanding the on-line availability of situation knowledge.

Fourth, this research contributes to theories of thematic roles. There is a logical lineage in the manner in which thematic roles have been considered over the years. Chomsky (1965) incorporated the notion of general selectional restrictions as part of the grammar. Katz and Fodor (1964) used the same basic notions, but considered selectional restrictions as part of lexical semantics. Both accounts incorporated a rather limited set of thematic roles and selectional features. Then two moves occurred. First it became apparent that a rich set of thematic roles was required, leading to an ever-expanding taxonomy. Second, researchers such as Johnson-Laird (1983) and McCawley (1968) pointed out that selectional restrictions must make reference to world knowledge and that some of this knowledge was verb-specific. More recently, Dowty (1991) proposed notions such as proto-agent and proto-patient features. This is an improvement because the

combinatorics generate a larger number of possible thematic roles. Furthermore, it also avoids the necessity of definitional or binary classification. Next, McRae and colleagues (Ferretti et al., in press; McRae et al., 1997; McRae et al., 1998) demonstrated that although these proposals captured important general facts that cut across verbs, they do not go far enough. Their research demonstrated that a larger and more detailed set of features are necessary to capture the empirical facts of language comprehension and production. The current research provides further evidence for these claims and extends them by demonstrating that the grammatical category of aspect constrains the activation of detailed thematic role knowledge during on-line thematic role assignment.

In conclusion, the research reported herein demonstrates the value of integrating various areas including theories of concepts, conceptual combination, thematic roles, inflectional morphology, and computational modeling. Integrating these areas has led to important insights regarding the processes that underlie language comprehension.

References

Altmann, G. T. M., Garnham, A., & Dennis, Y. (1992). Avoiding the garden path: Eye movements in context. *Journal of Memory and Language, 31*, 685-712.

Altmann, G. T. M., & Kamide, Y. (1999). Incremental interpretation at verbs: restricting the domain of subsequent reference. *Cognition, 73*, 247-264.

Anderson, S. J., & Conway, M. A. (1997). Representations of autobiographical memories. In M. A. Conway (Ed.), *Cognitive Models of Memory* (pp. 217-246). Cambridge, MA: MIT Press.

Barsalou, L. W. (1988). The content and organization of autobiographical memories. In U. Neisser and E. Winograd (Eds.), *Remembering reconsidered: Ecological and traditional approaches to the study of memory* (pp. 193-243). New York, NY: Cambridge University Press.

Bower, G. H., Black, J. B., & Turner, T. J. (1979). Scripts in memory for text. *Cognitive Psychology, 11*, 177-220.

Brown, N. R., & Schopflocher, D. (1998). Event clusters: An organization of personal events in autobiographical memory. *Psychological Science, 9*, 470-475.

Caplan, D., Hildebrandt, N., & Waters, G. S. (1994). Interaction of verb selectional restrictions, noun animacy, and syntactic form in sentence processing. *Language and Cognitive Processes, 9*, 549-585.

Carlson, G. N., & Tanenhaus, M. K. (1988). Thematic roles and language comprehension. In W. Wilkins (Ed.), *Thematic relations* (pp. 263- 288). New York: Academic Press.

Carrieras, M., Carriedo, N., Alonso, M.A, & Fernandez, Al. (1997). The role of verb tense and verb aspect in the foregrounding of information during reading. *Memory & Cognition*, *25*, 438-446.

Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.

Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments & Computers*, *25*, 257-271.

Cohen, B., & Murphy, G. L. (1984). Models of concepts. *Cognitive Science*, *8*, 27-58.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*, 407-428.

Collins, A.M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, *8*, 240-247.

Comrie, B. (1976). *Aspect: an introduction to the study of verbal aspect and related problems*. New York, NY: Cambridge University Press.

Dosher, B. A., & Corbett, A. T. (1982). Instrument inferences and verb schemata. *Journal of Verbal learning and Verbal Behavior*, *10*, 531-539.

Dowty, D. (1991). Thematic proto-roles and argument selection. *Language*, *67*, 547-619.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*, 179-211.

Ferretti, T. R., & McRae, K. (1999). Modeling the role of verb-bias and plausibility in the direct object / sentence complement ambiguity. In *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society* (pp. 161-166). Hillsdale NJ: Erlbaum.

Ferretti, T. R., McRae, K., & Hatherell, A. (in press). Integrating verbs, situation schemas and thematic role concepts. *Journal of Memory and Language*.

Fillmore, C. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), *Universals in linguistic theory* (pp. 1-90). New York: Holt, Rinehart, & Winston.

Fletcher, C. R. & Bloom, C. P. (1988). Causal reasoning in the comprehension of simple narrative texts. *Journal of Memory and Language*, *27*, 235-244.

Gagné, C. L. (2000). Relation-Based Combinations Versus Property-Based Combinations: A Test of the CARIN Theory and the Dual-Process Theory of Conceptual Combination. *Journal of Memory and Language*, *42*, 365-389.

Gagné, C. L., & Murphy, G. L. (1996). Influence of discourse context on feature availability in conceptual combination. *Discourse Processes*, *22*, 79-101.

Gagné, C. L., & Shoben, E.J. (1997). Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 71-87.

Gentner, D. (1975). Evidence for the psychological reality of semantic components: The verbs of possession. In D. A. Norman and D. E. Rumelhart (Eds.), *Explorations in Cognition* (pp. 211-246). San Francisco: W. H. Freeman.

Gentner, D. (1981). Some interesting differences between nouns and verbs. *Cognition and Brain Theory*, *4*, 161-178.

Goldberg, A. E. (1998). Patterns of experience in patterns of language. In M. Tomasello (Ed.), *The new psychology of language; Cognitive and*

functional approaches to language structure (pp. 203-219). Mahwah, NJ: Erlbaum.

Golden, R. M., & Rumelhart, D. E (1993). *A Parallel Distributed Processing Model of Story Comprehension and Recall. Discourse Processes*, *16*, 203-237.

Graesser, A. C., Woll, S. B., Kowalski, D. J., & Smith, D. A. (1980). Memory for typical and atypical actions in scripted activities. *Journal of Experimental Psychology: Human Learning and Memory*, *6*, 503-515.

Haspelmath, M. (1994). Passive Participles across Languages. In B. Fox and P. J. Hopper (Eds.). *Voice: form and function*. Amsterdam; Philadelphia: J. B. Benjamins.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, *93*, 411-428.

Jackendoff, R. (1983). *Semantics and cognition*. Cambridge, MA: MIT Press.

Johnson-Laird, P. N. (1983). *Mental Models*. Cambridge, England: Cambridge University Press.

Katz, J. J., & Fodor, J. A. (1964). The structure of a semantic theory. In J. A. Fodor and J. J. Katz (Eds.), *The structure of language* (pp. 479-518). Englewood Cliffs, NJ: Prentice Hall.

Kay, P. & Zimmer, K. (1976). On the semantics of compounds and genitives in English. *Sixth California Linguistics Proceedings*, 29-35.

Keefe, D. E., & McDaniel, M. A. (1993). The time course and durability of predictive inferences. *Journal of Memory and Language*, *32*, 446-463.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review*, *95*, 163-182.

Kolodner, J. L. (1983). Maintaining organization in dynamic long-term memory. *Cognitive Science*, *7*, 243-280.

Lancaster, J. L., & Barsalou, L. W. (1997). Multiple organisations of events in memory. *Memory*, *5*, 569-599.

Langacker, R. W. (1987). *Foundations of cognitive grammar, Vol. 1: Theoretical Prerequisites*. Stanford, CA: Stanford University Press.

Levi, J. N. (1978). *The syntax and semantics of complex nominals*. New York: Academic Press.

Lucas, M. M., Tanenhaus, M. K., & Carlson, G. N. (1990). Levels of Representation in the Interpretation of Anaphoric Reference and Instrument Inference. *Memory & Cognition*, *18*, 611-631.

Lyons, J. (1977). *Semantics*, *Vol 2*. Cambridge University Press.

MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review*, *101*, 676-703.

MacDonald, M. C. (1994). Probabilistic constraints on syntactic ambiguity resolution. *Language and Cognitive Processes*, *9*, 157-201.

MacWhinney, B. (1999). The emergence of language from embodiment. In B. MacWhinney (Ed.), *The emergence of language* (pp.213-256). Mahwah, N.J: Erlbaum.

Marslen-Wilson, W. D. (1975). Sentence perception as an interactive parallel process. *Science*, *189*, 226-228.

McCawley, J. D. (1968). The role of semantics in a grammar. In E. Bach & R. T. Harms (Eds.), *Universals in linguistic theory*. New York: Holt, Rinehart, & Winston.

McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General, 114*, 159-188.

McKoon, G., & Ratcliff, R. (1981). The comprehension processes and memory structures involved in instrumental inference. *Journal of Verbal learning and Verbal Behavior, 20*, 671-682.

McKoon, G., & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*, 82-91.

McKoon, G., Ratcliff, R., & Dell, G. S. (1986). *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*, 295-306.

McRae, K., Ferretti, T. R., & Amyote, L. (1997). Thematic roles as verb-specific concepts. *Language and Cognitive Processes, 12*, 137-176.

McRae, K., Spivey-Knowlton, M. J., & Tanenhaus, M. K. (1998). Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language, 38*, 283-312.

Moens, M., & Steedman, M. J. (1988). Temporal ontology and temporal reference. *Computational Linguistics, 14*, 15-28.

Minsky, M. (1975). A framework for representing knowledge. In P.H Winston (Ed.), *The psychology of computer vision* (pp. 211-277).

Morrow, D. G. (1985a). Prepositions and verb aspect in narrative understanding. *Journal of Memory and Language, 24*, 390-404.

Morrow, D. G. (1985b). Prominent Characters and Events Organize Narrative Understanding. *Journal of Memory and Language, 24*, 304-319.

Morrow, D. G. (1990). Spatial Models, Prepositions, and Verb-Aspect Markers. *Discourse Processes, 13*, 441-469.

Morrow, D. G., Greenspan, S. E., & Bower, G.H. (1987). Accessibility and situation models in narrative comprehension. *Journal of Memory and Language*, *26*, 165-187.

Moss, H. E., Ostrin, R. K., Tyler, L. K., & Marslen-Wilson, W. D. (1995). Accessing different types of lexical semantic information: Evidence from priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 863-883.

Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science*, *12*, 529-562.

Murphy, G. (1990). Noun phrase interpretation and conceptual combination. *Journal of Memory and Language*, *29*, 259-288.

Nakhimovsky, A. (1988). Aspect, aspectual class, and the temporal structure of narrative. Moens, M., & Steedman, M. J. (1988). Temporal ontology and temporal reference. *Computational Linguistics*, *14*, 15-28.

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective view of current findings and theories. In D. Besner & G. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264-336). Hillsdale, NJ: Lawrence Erlbaum.

Pichert, J. W., & Anderson, R. C. (1977). Taking different perspectives on a story. *Journal of Educational Psychology*, *69*, 309-315.

Pollatsek, A., & Well, A. D. (1995). On the use of counterbalanced designs in cognitive research: A suggestion for a better and more powerful analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 785-794.

Potter, M. C., & Faulconer, B. A. (1979). Understanding Noun Phrases. *Journal of Verbal Learning and Verbal Behavior*, *18*, 509-521.

Potts, G. R., Keenan, J. M., & Golding, J. M. (1988). Accessing the occurrence of elaborative inferences: Lexical decision versus naming. *Journal of Memory of Language*, *27*, 399-415.

Reiser, B. J., Black, J. B., & Abelson, R. P. (1985). Knowledge structures in the organization and retrieval of autobiographical memories. Cognitive *Psychology*, *17*, 89-137.

Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573-605.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382-439.

Rumelhart, D. E. (1980). Schemata: The building blocks of cognition. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), *Theoretical issues in reading comprehension* (pp. 33-58). Hillsdale, NJ: Erlbaum.

Rumelhart, D. E. (1991). Understanding Understanding. In W. Kessen, A. Ortony, & F. Craik (Eds.), *Memories, Thoughts, and Emotions: Essays in Honor of George Mandler* (pp. 257-275). Hillsdale, NJ: Erlbaum.

Rumelhart, D. E., & Levin, A. (1975). A Language Comprehension System. In D. A. Norman and D. E. Rumelhart (Eds.), *Explorations in Cognition* (pp. 179-208). San Francisco: W. H. Freeman.

Rumelhart, D. E., & Ortony, A. (1977). The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge* (pp. 99-135). Hillsdale, NJ: Erlbaum.

Rumelhart, D. E., Smolensky, J.L., McClelland, and Hinton, G. E. (1986). Schemata and Sequential Thought Processes in PDP models. *In Parallel*

Distributed Processing: Explorations in the Microstructure of Cognition, Vol 2, (pp. 7-57). Cambridge, MA: MIT Press.

Sanford, A. J., & Garrod, S. C. (1981). *Understanding written language.* Chichester: John Wiley.

Schlesinger, I. M. (1995). *Cognitive space and linguistic case.* Cambridge: Cambridge University Press.

Schmauder, A. R., & Egan, M. C. (1998). The influence of semantic fit on on-line sentence processing. *Memory and Cognition, 26*, 965-978.

Seidenberg, M. (1992). Connectionism without tears. In S. Davis (Ed.), *Connectionism: Theory and practice. Vancouver studies in cognitive science*, Vol. 3 (pp. 84-122). Oxford, England: Cambridge University Press.

Shank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding.* Hillsdale, NJ: Erlbaum.

Sharkey, J. C., & Sharkey, N. E. (1992). Weak Contextual Constraints in Text and Word Priming. *Journal of Memory and Language, 31*, 543-572.

Shelton, J. R., & Martin, R. C. (1992). How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18*, 1191-1210.

Smith, C. S. (1991). *The Parameter of Aspect.* Kluwer Academic Publishers: Boston.

Smith, E. D. & Osherson, D. N., Rips, L. J., & Keane, M (1988). Combining prototypes: A modification model. *Cognitive Science, 12*, 485-527.

Spivey-Knowlton, M. J. (1996). *Integration of visual and linguistic information: Human data and model simulations.* Unpublished doctoral dissertation, University of Rochester, Rochester, N.Y.

Spivey-Knowlton, M. J., & Sedivy, J. C. (1995). Resolving attachment ambiguities with multiple constraints. *Cognition*, *55*, 227-267.

Spivey, M. J., & Tanenhaus, M. K. (1998). Syntactic ambiguity resolution in discourse: Modeling the effects of referential context and lexical frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 1521-1543.

Springer, K., & Murphy, G. L. (1992). Feature availability in conceptual combination. *Psychological Science*, *3*, 111-117.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension, *Science*, *268*, 1632-1634.

Taraban, R., & McClelland, J. L. (1988). Constituent attachment and thematic role assignment in sentence processing: Influences of content-based expectations. *Journal of Memory and Language*, *27*, 597-632.

Till, R. E., Mross, E. F., & Kintsch, W. (1988). Time course of priming for associate and inference words in a discourse context. *Memory and Cognition*, *16*, 283-298.

Trueswell, J., Tanenhaus, M., & Kello, C. (1993). Verb specific constraints in sentence processing: Separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 528-553.

Truitt, T. P., & Zwaan, R. A. (1997). Verb aspect affects the generation of instrumental inferences. *Paper presented at the 38th Annual Meeting of the Psychonomic Society, Philadelphia, November, 1997.*

Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford University Press.

Tversky, B., & Hemenway, K. (1983). Categories of environmental scenes. *Cognitive Psychology, 15*, 121-149.

van Dijk, T. A., & Kintsch, W. (1983). *Strategies of discourse comprehension.* New York: Academic Press.

Vu, H., Kellas, G., & Paul, S. T. (1998). Sources of constraint on lexical ambiguity resolution. *Memory and Cognition, 26*, 979-1001.

Wisniewski, E. J. (1996). Construal and similarity in conceptual combination. *Journal of Memory and Language, 35*, 424-453.

Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin and Review, 4*, 167-183.

Wisniewski, E. J., & Gentner, D. (1991). On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In G. B. Simpson (Ed.), *Understanding word and sentence*, pp. 241-284). Amsterdam: North Holland.

Wisniewski, E. J., & Love, B. C. (1998). Relations versus properties in conceptual combination: *Journal of Memory and Language, 38*, 177-202.

Zwaan, R. A., Langston, M. C., & Graesser, A. C. (1995). The construction of situation models in narrative comprehension: An Event-Indexing Model. *Psychological Science, 6*, 292-297.

Appendix A

Experiment 1 related verb-location and instrument pairs with their accompanying ratings. Note that the verbs appeared in both their past imperfective form (was verb-ing) and past perfect form (had verb-ed). Irregular past participle forms are indicated in parentheses.

Related Verb	*Location*	*Rating*	*Unrelated Verb*
cook	kitchen	7.0	worship
exercise	gymnasium	7.0	mourn
drive (had driven)	highway	7.0	preach
skate	arena	7.0	pray
sleep (had slept)	bedroom	7.0	browse
dance	ballroom	7.0	examine
eat (had eaten)	restaurant	7.0	sleep
hunt	forest	6.9	study
act	theater	6.9	bury
study	library	6.9	skate
shower	bathroom	6.9	sketch
gamble	racetrack	6.8	swam
pray	temple	6.8	drove
mourn	cemetery	6.8	act
stroll	park	6.8	confess
preach	synagogue	6.7	exercise
applaud	stadium	6.7	trawl
worship	chapel	6.7	stroll
browse	museum	6.6	cook
swim (had swam)	ocean	6.6	bet

draw (had drawn)	studio	6.5	shower
bury	sandbox	6.3	wait
fish	river	6.2	track
confess	court	6.0	dance

Related Verb	*Instrument*	*Rating*	*Unrelated Verb*
stir	spoon	6.9	shoot
slap	hand	6.8	devour
paint	brush	6.8	teach
dust	rag	6.8	cut
heat	furnace	6.7	watch
eat (had eaten)	fork	6.7	hit
hunt	rifle	6.7	trawl
repair	hammer	6.7	track
drink (had drunk)	cup	6.6	slap
shoot (had shot)	pistol	6.5	stir
fish	bait	6.4	paint
wash	detergent	6.3	revive
teach (had taught)	textbook	6.2	wash
gamble	dice	6.2	vandalize
watch	binoculars	6.1	heat
revive	oxygen	6.0	repair
stab	dagger	5.9	float
cut (had cut)	saw	5.9	dust
hit (had hit)	bat	5.9	serve
wrap	tinfoil	5.9	bet
inflate	pump	5.8	dug

serve	platter	5.8	sketch
draw (had drawn)	marker	5.8	drink
dig (had dug)	spade	5.6	inflated

Appendix B

Experiment 2 and 4 participle-agent and participle-patient pairs and their agenthood (Ahood) and patienthood (Phood) ratings. Note that the participles appeared in both their present (participle+ing) and past (participle+ed) forms. Irregular past participles are indicated in parentheses. Note the asterisks indicate the two agent items in Experiment 4 that did not appear in Experiment 2.

Participle	*Agent*	*Ahood*	*Phood*
dismiss	principal	6.3	2.0
kiss	lover	6.9	6.6
invite	host	6.5	2.2
instruct	coach	6.7	2.7
love	grandmother	6.6	6.6
sentence	judge	6.9	1.3
attack	troop	6.6	6.2
lift	worker	6.1	2.1
investigate	detective	6.4	1.9
carry	postman	6.8	1.6
terrorize	pirate	6.5	2.2
approve	committee	6.2	3.8
slaughter	butcher	6.2	1.4
pay (paid)	customer	6.5	1.6
grad	teacher	6.9	2.6
accuse	prosecutor	6.8	2.6
teach (taught)	professor	6.6	2.6
interview	reporter	6.8	2.7

capture	policeman	6.4	2.1
scrub	janitor	6.8	1.4
request	specialist	5.7	5.5
beat (beaten)	bully	6.5	1.9
search	patrolman	6.3	1.4
choose (chosen)	contestant	6.0	5.3
arrest	cop	6.7	1.6
read (read)	philosopher	6.5	3.0
serve	waitress	6.8	2.5
scratch	cat	6.9	4.1
record	singer	6.7	5.5
recognize	witness	6.1	3.9
shoot (shot)	hunter	6.9	2.8
throw (thrown)	pitcher	6.8	1.6
audit	accountant	6.0	4.0
torture	rapist	6.6	2.2
write (written)	author	6.7	1.8
visit	tourist	6.8	1.4
sketch	artist	6.5	2.5
show (shown)	salesman	6.6	1.6
stalk	prowler	6.4	2.1
study	scientist	6.6	2.4
stole (stolen)	criminal	6.7	1.4
question	lawyer	6.5	2.9
find (found)	archeologist	6.4	1.8
fire	employer	6.1	2.4
frighten	monster	6.4	1.8

devour	snake	6.2	3.9
govern	president	6.7	2.7
fight (fought)	boxer	6.8	6.1
disobey	brat	6.7	3.1
adopt	parent	5.8	1.7
entertain	comedian	6.9	2.7
execute	terrorist	6.1	4.0
evaluate	instructor	6.1	4.0
kick	donkey	6.1	3.4
convict	juror	6.6	1.3
accept	friend	6.1	5.8
cheer	spectator	6.7	1.3
conquer	king	6.1	3.8
cure	doctor	6.8	3.8
hire	boss	6.7	2.9
help	fireman	6.7	3.9
admire	fan	6.8	1.7
describe	person	6.5	6.3
interrogate	inspector	6.3	1.6
*identify	nurse	5.3	4.2
*eat (eaten)	man	6.9	1.7

Participle	*Patient*	*Ahood*	*Phood*
invite	guest	1.9	6.4
torture	slave	1.3	5.6
startle	bird	3.2	6.0
draw (drawn)	model	2.0	6.4
eat (eaten)	chicken	5.9	6.7
transport	cattle	2.0	5.6
question	witness	2.0	6.7
rescue	hostage	1.4	5.2
admire	athlete	4.8	6.4
entertain	audience	1.7	6.7
overthrow (overthrown)	dictator	4.3	5.3
search	thief	5.7	6.1
adopt	baby	1.4	6.4
present	nominee	4.2	5.5
visit	family	6.4	6.6
evaluate	applicant	3.3	5.7
carry	newborn	1.2	6.2
attack	infantry	5.6	5.7
shoot (shot)	deer	1.0	6.4
hunt	rabbit	2.4	6.7
recognize	celebrity	3.8	6.8
grade	student	2.4	6.8
worship	king	2.6	6.2
kick	wimp	1.6	5.4
investigate	suspect	2.4	6.4
sketch	woman	4.3	5.3

interrogate	culprit	2.0	5.4
release	inmate	1.5	5.9
fire	employee	1.9	6.4
accuse	defendant	3.6	6.8
lift	infant	2.5	5.9
cure	patient	1.4	6.1
love	husband	6.6	6.8
choose (chosen)	candidate	4.0	5.6
applaud	musician	2.9	6.8
arrest	crook	1.2	5.9
dismiss	pupil	1.5	5.5
convict	criminal	1.4	5.9
terrorize	victim	1.4	6.6
frighten	cat	3.1	5.5
serve	customer	1.5	7.0
scratch	gardener	3.4	5.7
slaughter	pig	1.0	6.8
lecture	freshman	1.9	5.4
govern	villager	3.2	6.5
punish	child	1.5	5.8
instruct	novice	1.6	5.3
audit	taxpayer	2.7	6.2
cheer	quarterback	3.8	5.2
kiss	princess	5.9	5.9
enslave	peasant	1.7	5.6
stalk	prey	1.7	6.0
capture	fugitive	2.6	5.1

teach (taught)	trainee	1.6	6.3
pay (paid)	cashier	3.1	5.4
chase	mouse	3.1	5.5

Appendix C

Complete example of the model processing the phrase governing villager over time (i.e., across cycles of competition).

Cycle	*Δcrit*	*Agent Act*	*Patient Act*	*Ahood*	*Phood*	*Ing Bias*	*Ed Bias*
1	.993	.49	.51	.29	.71	.69	.31
2	.986	.49	.51	.28	.72	.69	.31
3	.979	.48	.52	.28	.72	.69	.31
4	.972	.48	.52	.28	.72	.68	.32
5	.965	.48	.52	.28	.72	.68	.32
6	.958	.47	.53	.27	.73	.68	.32
7	.951	.47	.53	.27	.73	.67	.33
8	.944	.46	.54	.26	.74	.67	.33
9	.937	.46	.54	.26	.74	.66	.34
10	.930	.45	.55	.25	.75	.65	.35
11	.923	.44	.56	.24	.76	.64	.36
12	.916	.43	.57	.24	.76	.63	.37
13	.909	.42	.58	.23	.77	.62	.38
14	.902	.41	.59	.22	.78	.60	.40
15	.895	.40	.60	.20	.80	.59	.41
16	.888	.38	.62	.19	.81	.57	.43
17	.881	.36	.64	.18	.82	.54	.46
18	.874	.34	.66	.16	.84	.52	.48
19	.867	.31	.69	.14	.86	.48	.52
20	.860	.29	.71	.13	.87	.45	.55
21	.853	.26	.74	.11	.89	.40	.60
22	.846	.22	.78	.09	.91	.36	.64
23	.839	.19	.81	.07	.93	.31	.69
24	.832	.16	.84	.06	.94	.26	.74

Appendix D

Experiment 3 participle-agent and participle-patient pairs and their agenthood (Ahood) and patienthood (Phood) ratings. Note that the participles appeared in both their present (participle+ing) and past (participle+ed) forms. Irregular past participles are indicated in parentheses.

Participle	*Agent*	*Ahood*	*Phood*
fire	employer	6.1	2.4
punish	parent	6.5	1.5
devour	snake	6.2	3.9
accuse	prosecutor	6.8	2.6
cheer	spectator	6.7	1.3
identify	nurse	5.8	4.2
show (shown)	salesman	6.6	1.6
frighten	monster	6.4	1.8
carry	postman	6.8	1.6
fire	owner	6.3	1.9
help	fireman	6.7	3.9
visit	tourist	6.8	1.4
love	grandmother	6.6	6.6
torture	rapist	6.6	2.2
sentence	judge	6.9	1.3
conquer	king	6.1	3.8
approve	committee	6.2	3.8
search	patrolman	6.3	1.4
interrogate	inspector	6.3	1.6

dismiss	principal	6.3	2.0
describe	person	6.5	6.3
disobey	brat	6.7	3.1
stalk	prowler	6.4	2.1
lift	workers	6.1	2.1
entertain	comedian	6.9	2.7
teach (taught)	professor	6.6	2.6
request	specialist	5.7	5.5
interview	reporter	6.8	2.7
study	scientist	6.6	2.4
evaluate	instructor	6.1	4.0
terrorize	pirates	6.5	2.2
serve	butler	6.7	2.1
eat (eaten)	man	6.9	1.7
slaughter	butcher	6.2	1.4
consider	scientists	6.6	4.0
kick	donkey	6.1	3.4
shoot (shot)	hunter	6.9	2.8
cure	doctor	6.8	3.8
choose (chosen)	contestant	6.0	5.3
find (found)	archeologist	6.4	1.8
scrub	janitor	6.8	1.4
audit	accountant	6.0	4.0
convict	juror	6.6	1.3
beat (beaten)	bully	6.5	1.9
arrest	cop	6.7	1.6
write (written)	author	6.7	1.8

execute	terrorists	6.1	4.0
admire	fan	6.8	1.7
steal (stolen)	criminal	6.7	1.4
question	lawyer	6.5	2.9
grade	teacher	6.9	2.6
read (read)	philosopher	6.5	3.0
hire	boss	6.7	2.9
throw (thrown)	pitcher	6.8	1.6
govern	president	6.7	2.7
pay (paid)	customer	6.5	1.6
attack	troops	6.6	6.2
accept	friend	6.1	5.8
record	singer	6.7	5.5
kiss	lover	6.9	6.6
invite	host	6.5	2.2
sketch	artist	6.5	2.5
recognize	witness	6.1	3.9
instruct	coach	6.7	2.7
scratch	cat	6.9	4.1
fight (fought)	boxer	6.8	6.1
capture	policeman	6.4	2.1
serve	waitress	6.8	2.5
convict	judge	6.4	1.5
adopt	parent	5.8	1.7
shoot (shot)	assassin	6.9	2.2
investigate	detective	6.4	1.9

Participle	*Patient*	*Ahood*	*Phood*
kiss	princess	5.9	5.9
audit	taxpayer	2.7	6.2
recognize	celebrity	3.8	6.8
cure	patient	1.4	6.1
stalk	prey	1.7	6.0
applaud	musician	2.9	6.8
eat (eaten)	chicken	5.9	6.7
frighten	cat	3.1	5.5
sketch	woman	4.3	5.3
scratch	gardener	3.4	5.7
adopt	baby	1.4	6.4
arrest	crook	1.2	5.9
love	husband	6.6	6.8
fire	employee	1.9	6.4
search	thief	5.7	6.1
sentence	criminal	1.3	6.7
enslave	peasants	1.7	5.6
cheer	quarterback	3.8	5.2
carry	newborn	1.2	6.2
investigate	suspect	2.4	6.4
admire	athlete	4.8	6.4
capture	fugitive	2.6	5.1
torture	slave	1.3	5.6
dismiss	pupil	1.5	5.5
entertain	audience	1.7	6.7
grade	student	2.4	6.8

devour	rabbit	3.3	5.7
startle	bird	3.2	6.0
chase	mouse	3.1	5.5
interrogate	culprit	2.0	5.4
draw (drawn)	model	2.0	6.4
present	nominee	4.2	5.5
worship	king	2.6	6.2
punish	child	1.5	5.8
rescue	hostage	1.4	5.2
terrorize	victim	1.4	6.6
serve	customer	1.5	7.0
transport	cattle	2.0	5.6
shoot (shot)	deer	1.0	6.4
release	inmate	1.5	5.9
evaluate	applicant	3.3	5.7
hunt	rabbit	2.4	6.7
investigate	gangster	3.5	6.4
accuse	defendant	3.6	6.8
kick	wimp	1.6	5.4
attack	infantry	5.6	5.7
lift	infant	2.5	5.9
chose (chosen)	candidate	4.0	5.6
visit	family	6.4	6.6
lecture	freshman	1.9	5.4
slaughter	pig	1.0	6.8
question	witness	2.0	6.7
invite	guest	1.9	6.4

teach (taught)	trainee	1.6	6.3
interview	applicant	1.6	6.6
govern	villager	3.2	6.5
convict	criminal	1.4	5.9
pay (paid)	cashier	3.1	5.4
overthrow (overthrown)	dictator	4.3	5.3
instruct	novice	1.6	5.3