ELSEVIER

Journal of Cognitive Systems Research 1 (1999) 19–33

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

Not propositions

Action editor: Gregg C. Oden

Arthur M. Glenberg*, David A. Robertson, Jennifer L. Jansen, Mina C. Johnson-Glenberg

Department of Psychology, University of Wisconsin-Madison, 1202 West Johnson Street, Madison, WI 53706, USA

Received 4 May 1999; received in revised form 14 July 1999; accepted 20 July 1999

Abstract

Current computational accounts of meaning in the cognitive sciences are based on abstract, amodal symbols (e.g., nodes, links, propositions) that are arbitrarily related to their referents. We argue that such accounts lack convincing empirical support and that they do not provide a satisfactory account for linguistic meaning. One historic set of results supporting the abstract symbol view has come from investigation into comprehension of negated sentences, such as ‘‘The buttons are not black.’’ These sentences are presumed to be understood as two propositions composed of abstract symbols. One proposition corresponds to ‘‘the buttons are black,’’ and it is embedded in another proposition corresponding to ‘‘it is not true.’’ Thus, the propositional account predicts (a) that comprehension of negated sentences should take longer than comprehension of the corresponding positive sentence (because of the time needed to construct the embedding), but (b) that the resulting embedded propositions are informationally equivalent (but of opposite valence) to the simple proposition underlying the positive sentence. Contrary to these predictions, Experiment 1 demonstrates that negated sentences out of context are interpreted as situationally ambiguous, that is, as conveying less specific information than positive sentences. Furthermore, Experiment 2 demonstrates that when negated sentences are used in an appropriate context, readers do not take longer to understand them. Thus, difficulty with negation is demonstrated to be an artifact of presentation out of context. After discussing other serious problems with the use of abstract symbols, we describe the Indexical Hypothesis. This embodied account of meaning does not depend on abstract symbols, and hence it provides a more satisfactory account of meaning.

1. Not propositions

What does ‘‘not propositions’’ mean? Is this article about negated propositions? Does it assert that propositions do not apply in some domain? And if not propositions, then what? In fact, this article is about these questions and about the very phenomenon illustrated by the questions: out of context, negated sentences are often more ambiguous than the corresponding positives. However, as we will demonstrate, when negated sentences are used in appropriate contexts, all difficulty disappears. This simple effect has major implications: Because the difficulty with negatives has been used as a key support for the theory of propositions and abstract symbols, removing that support suggests that it is time to question the hegemony of propositional theory, abstract symbols, and the presumption that cognition is a computational phenomenon. Instead, we will argue for a

*Corresponding author. Tel.: 608-262-8992; fax: 608-262-4029.
E-mail address: glenberg@facstaff.wisc.edu (A.M. Glenberg)

1389-0417/99/$ – see front matter
PII: S1389-0417(99)00004-2

theory of meaning based on embodiment (e.g., Glenberg, 1997; Lakoff, 1987).

We begin with a discussion of how psychologists have used propositions as a theory of meaning. Then, we review some of the empirical evidence that has been used to support the theory, in particular work on negated sentences. Next, we present the results from three experiments that undermine that support. Finally, we discuss general problems with propositional theory and present an alternative account of meaning. This alternative is based on the idea of embodiment (e.g., Barsalou, 1993; Glenberg, 1997; Lakoff, 1987), that cognition is intimately connected with the functioning of the body. The Indexical Hypothesis (Glenberg & Robertson, in press (a); Glenberg & Robertson, in press, (b), see also Robertson et al., 1999) provides a link between the general theory of embodied meaning and its specific application to language.

2. Propositions

According to most psychological accounts of language comprehension, meaning involves deriving propositions from discourse and connecting them appropriately (e.g., Carpenter & Just, 1992; Kintsch, 1988). Often, the appropriate connection involves adding and inferring propositions that may be generalizations (e.g., macro-propositions in Van Dijk & Kintsch, 1983), or explanations (Van den Broek, 1994), or mental models (Van Dijk & Kintsch, 1983). All of these efforts are based on the central idea of a proposition. If that idea is wrong, then there is good reason to question these accounts.

Psychologists commonly use the term ''proposition'' as a shorthand for ''mental representation of an assertion as an expression in a propositional calculus.'' In these psychological theories, the proposition typically consists of a relational term and its arguments. Thus, the fact that the printer is next to the monitor might be propositionalized as ''next-to (printer, monitor),'' and the fact that the monitor is on the table might be propositionalized as ''on (monitor, table).'' In these examples, the relations are ''next-to'' and ''on'' and the arguments are ''printer,'' ''monitor,'' and ''table.'' Some theories allow for unary relations to describe properties of objects such as color, and *n*-ary relations to describe more complex ideas such as giving, in which there may be an actor, recipient, object, mode, time, place, and so on. Propositions are thought to be the smallest units to have ''truth values,'' that is, about which it can be asserted that they are true or false. In all propositional theories, language comprehension entails forming propositions based on the words in sentences. Section 3.4 lays out a variety of problems with this view of language understanding, as well as an alternative, the Indexical Hypothesis. For now, we turn to some of the evidence that has been offered in support of propositional theories.

There are several good reasons for the predominance of propositional accounts of cognition. Propositional theories work very well in formal domains such as logic, mathematics, and computer sciences. There are also historical reasons for the predominance of propositional theories. At the time when propositional theories were being developed, it seemed clear that simple associative theories could not handle the complexities of cognition, and there were many arguments offered against systems based on simple analogical representations such as visual images (e.g., Pylyshyn, 1973). Finally, three empirical phenomena have been offered as grounds for supposing that language understanding involved the construction of propositions. These three are that recall of text often seems to reflect propositional units and gist (Kintsch, 1974), priming seems to follow propositional relations rather than contiguity (McKoon & Ratcliff, 1980), and the difficulty in processing negated sentences. We think that all of the empirical results can be interpreted without recourse to propositions, but for this paper we concentrate on the results from negated sentences.

In many propositional theories, the representation of a negated sentence is more complex than that of the corresponding positive. The complexity arises because of the necessity of representing both the positive relation and its negation. Thus the sentence ''The monitor is not on the table'' might be represented by ''not (on (monitor, table)).'' This complexity ought to have cognitive implications (e.g., slowing processing), and this is just what a series of experiments using the sentence–picture verification task demonstrated (Carpenter & Just, 1975; Clark & Chase, 1972; Trabasso et al., 1971). Using this procedure, the participant reads a simple sentence (e.g., ''The star is above the cross'') and views a picture illustrating a relation between a star and a cross. The task is to respond affirmatively if the sentence is a correct description of the picture. A basic finding from this procedure is that response to negative sentences (e.g., ''The star is not above the cross'') is slower than response to the affirmative. Also, the pattern of response time effects can be closely modeled by considering the number of comparisons between the presumed propositional representation of the sentence and the presumed propositional representation of the picture needed to verify that the representations are congruent. This close modeling was taken as strong support for the use of propositions as the basic format of representation.

The close fit notwithstanding, the sentence–picture verification procedure and the results have been criticized. Tanenhaus et al. (1976) note the following. First, the propositional representations presumed to be formed for the sentence–picture verification task are particular to the task: if they were used outside of the tight constraints of the experiment, they would generate many errors of reasoning. Second, as illustrated in Section 3.4, forming the appropriate propositional representation seems to require first understanding (whatever that might be), rather than understanding resulting from the propositions themselves. Third, the apparent close fit between response time and number of comparisons depends on postulating unmotivated rules for how the comparisons are done. Importantly, in order to account for the data, the rules have to be changed in response to the data themselves (e.g., Carpenter & Just, 1975), rather than to any features of the experiment.

There have also been several empirical demonstrations that the propositional account of negated sentences is, at best, incomplete. Wason (1965) attempted to demonstrate a ''context of plausible denial.'' That is, within the right pragmatic context, negated sentences are processed with felicity, contrary to propositional theory. According to Wason, the pragmatic context that supports negation is when an expectation needs to be denied, or when an exception needs to be noted. Unfortunately, Wason's methodology does not seem to meet current standards or rigor. In particular, in his experiment, the presentation of critical stimuli was preceded by the participant giving an uncontrolled description of the to-be-queried stimuli. Also, the reaction time analysis included both correct and incorrect responses.

De Villiers & Flusberg (1975) replicated Wason (1965) in a developmental study using a stronger methodology. Two-, three-, and four-year-old children were shown displays of seven similar objects (e.g., toy cars) and one dissimilar object (e.g., a baby bottle). The experimenter pointed to an object and prompted children with sentence fragments such as ''This is a . . . '' and ''This is not a . . . .'' The children responded with the name of the object, and the dependent variables were the time taken to name the object and the error rate. When the experimenter pointed to one of the seven similar objects, the children were faster in responding to the positive prompt than to the negated one. However, when the experimenter pointed to the dissimilar object (a plausible context for denial), the children were equally fast in responding to both the positive and negated prompts. Apparently, negated sentences are not uniformly difficult (even for two-year-old children), as the propositional account predicts.

More recent work on negation has been reported by MacDonald and Just (1989). They demonstrated that negation may suppress accessibility of terms. In their experiments, participants read sentences such as ''Almost every weekend, Elizabeth bakes some bread but no cookies for the children.'' Relative accessibility was demonstrated using a probe task (recognition and naming latencies). MacDonald and Just found that the negated term was less accessible than the non-negated term. Interestingly, however, they found little effect of negation on sentence reading time. They concluded that effects of negation depend on the level of representation being tapped, and that it is at the discourse level that negation reduces accessibility of concepts. We will have little else to say about the MacDonald and Just findings, because our experiments are not designed to test their claims.

In summary, there are logical (Tanenhaus et al., 1976), theoretical (Wason, 1965), and empirical (De Villiers & Flusberg, 1975) grounds for questioning the propositional account of negation and the sentence–picture verification procedure in particular.

Perhaps for these reasons, the sentence–picture verification procedure has dropped from favor. *Nonetheless, a conclusion based on this procedure remains embedded in our theories of cognition. That conclusion is that understanding is in terms of propositions, and understanding of negated sentences is in terms of embedded propositions.* If we can demonstrate that the propositional account of negatives is incorrect, that will remove one of the major pieces of empirical support for propositional theory. Unlike the situation in the 1960s and 1970s, there are now accounts of cognition, such as embodiment, which provide an alternative to propositions. Thus, we judged it valuable to revisit the question of understanding of negated sentences.

3. Preview of the experiments

Givon (1978) presents an analysis of negated sentences different from the propositional account, and from it we generated several contrasting predictions that form the basis of our experiments. We focus on two aspects of Givon's analysis. First, Givon notes that, out of context, negated sentences are in some ways more ambiguous than the corresponding positives, whereas that is not true of the supposed propositional counterparts. Thus, in a propositional theory, one can assert a proposition, P, or its negative, not-P, and the two can be used with equal certainty in propositional reasoning. Furthermore, not(not-P) is always equivalent to P. Now, consider language. The information content of ''The star is above the cross'' seems greater than ''The star is not above the cross.'' The former sentence seems to assert something definite: a definite location of the star. The latter sentence seems to give less information: where not to find the star, not where to find it. That is, out of context, the negated sentence is situationally ambiguous; it corresponds to a greater variety of situations than does the positive sentence. Furthermore, in language, some of the ideas consistent with ''The star is not above the cross'' are ''The start is next to the cross,'' ''The star is behind the cross,'' and ''There is no star,'' but negating any of these sentences does not bring us back to the original.

For the three experiments reported here, we restricted our investigation to negation explicitly marked with the word ''not''. There are other forms of negation including marked cases (with the prefix ''un-'' or the words ''no'', ''nothing'', etc.), unmarked cases (e.g., ''He lacks the necessary money,'' and ''The tree is short''), and pragmatic uses such as irony or sarcasm (Giora, 1995). Horn (1989) provides a more complete presentation of the aspects of negation. We restricted our investigations in the interest of experimental rigor. Also, much of the work with negated sentences that has been used to support propositional theory was conducted under a similar restriction.

Experiment 1 demonstrates the situational ambiguity of negated sentences using ratings. Participants were asked to rate the situational ambiguity of positive and negated sentences. We also varied whether the positive term was close to the end of a dimension (e.g., ''The car was stopped,'' is at the end of the speed dimension). We suspected that this would affect the ratings and would thereby demonstrate that the participants were not executing a mindless strategy of rating negated sentences as more ambiguous than positive sentences. Consistent with Givon's analysis, our results show that participants judge negated sentences out of context to be more situationally ambiguous than positive sentences. This finding speaks against propositional theory because negated and positive propositions should be equally informative.

The second part of Givon's analysis is similar to Wason's context of plausible denial. According to Givon, negated sentences seem to require a particular pragmatic context, and that within that context they are understood with facility. Consider a modification of Givon's example involving a conversation between two friends, A and B. A asks, ''What's new?'' B replies, ''I'm not pregnant.'' The reply seems distinctly odd, and might lead A to say, ''Huh? Oh, I didn't know that you were supposed to be pregnant.'' However, in a different context, B's reply is not odd and easily understood. Suppose that A and B had had a previous conversation (even months before) in which B had revealed that she was attempting to get pregnant. Then, B's reply would be easily interpreted by A in the light of that previous conversation. Givon proposes that felicitous use of negated sentences requires that they are used to counter presuppositions held by a listener/comprehender.

In Experiment 2 we contrast the reading times of negated and positive sentences when there is or is not a supporting context that provides the presuppositions. As Givon's analysis leads us to suspect, the difference in reading times between positive and negated sentences is virtually eliminated when there is a supporting context. That is, when negated sentences are used in pragmatically correct contexts, they are not more difficult to comprehend than affirmative sentences. Thus, the major empirical finding lending support to the propositional account of negated sentences (slow processing) only appears when pragmatic constraints are violated.

In Experiment 3 we demonstrate that our materials are not odd, and the experiment provides a connection to the sentence–picture verification procedure. We use the critical sentences from Experiments 1 and 2 in the sentence–picture task, and we demonstrate that out of context, these sentences produce results substantially identical to those found with the standard materials, namely much slower reading and verification times for the negated sentences than for their positive counterparts. Thus, we demonstrate that what has become the standard finding (slower processing of negated sentences) is actually an artifact of presentation out of context.

3.1. Experiment 1

Participants rated whether the sentences described a relatively specific situation or a range of situations. These sentences varied as to whether they were positive or negated and as to whether the sentences referred to the middle of a dimension or the end of a dimension.

3.1.1. Method

Participants. Thirty students in an introductory psychology course participated for extra credit. All participants were native English speakers.

Materials and design. Eight sentence sets were constructed to tap the dimensions of color, speed, temperature, lifespan, weight, capacity, volume, and contentment. Two of the sentences in each set were an end-dimension sentence (''The buttons are black'') and its negated form (''The buttons are not black''). The other two sentences were the mid-dimension sentence (''The buttons are darkly colored'') and its negated form (''The buttons are not darkly colored''). Half of the mid-dimension sentences were judged to refer to the same side of the continuum as the end dimension sentence and half were judged to refer to the other side. Because this variable did not appear to affect any of the results, it will not be mentioned again.

The sentences were presented on four pages. Each page contained eight sentences, one from each set. Next to each sentence was a rating scale consisting of the numbers one to six. No page had more than three exemplars of any of the four types of sentences. The order of the four pages was randomly varied across participants.

Procedure. Participants were instructed to rate the ambiguity of each sentence by circling a number from one (unambiguous) to six (ambiguous). Participants were instructed that an unambiguous sentence describes a specific situation and that an ambiguous sentence is consistent with many situations. To illustrate these definitions participants were told that the sentence ''The door is open'' is ambiguous because the door may be open a little or a lot. They were further instructed that the sentence ''The door is closed'' is less ambiguous because the sentence is consistent with fewer situations. Note that these instructions and examples did not used negated sentences so as not to bias the ratings of the negated sentences.

3.1.2. Results and discussion

The data of interest are presented in Table 1. All analyses were conducted with a Type I error rate set to .05.

Participants rated negated sentences (mean rating = 4.18) as more ambiguous than positive sentences (3.03), $F'(1, 31) = 21.46$, where F' is a quasi-F statistic from an analysis treating both

Table 1
Mean situational ambiguity ratings from Experiment 1

Sentence type	End-dimension	Mid-dimension
Positive	2.45	3.60
Negated	4.03	4.33

subjects and sentences as random factors (Clark 1973; Satterthwaite, 1946).

Thus, contrary to some propositional theories of language, negated sentences (and their corresponding propositions) are not informationally equivalent to the corresponding positive sentences. Also, mid-dimension sentences (3.96) were rated as more ambiguous than end-dimension sentences (3.24), $F'(1, 9) = 7.53$. Importantly, the effect of negation was much larger for end-dimension sentences than for mid-dimension sentences, $F'(1, 14) = 5.03$. Thus, participants seemed to be considering the meanings of the sentences and not just rating negated sentences as more ambiguous than the positive sentences.

We draw two conclusions from these results. The first conclusion is that the types of simple negated sentences used in the picture verification studies (e.g., Carpenter & Just, 1975) are more ambiguous than the positive counterparts. This ambiguity could be what is behind the increase in processing time, rather than the presumed complexity of the propositional representation.

The second conclusion is that these results challenge most propositional theories of language. Well-formed propositions, whether positive or negated, are equally informative. If language is understood by the direct translation of sentences into propositions, then the positive and negated sentences ought to have been interpreted as equally specific. Is this simple translation process the way propositional theories presume negation is understood? The question is difficult to answer, because many of the psychological accounts of language published since the late 1970s do not address negation explicitly.[1] One can imagine building a propositional theory to include a metric of information based on extensions of the propositions. For example, the metric might be the number or density of situations consistent with the proposition. Such a theory would seemingly account for the results of Experiment 1. Furthermore, the theory could account for situations in which negation apparently decreases ambiguity rather than increases ambiguity, such as ‘‘The door is (not) open.’’ The proposition ‘‘open(door)’’ is consistent with a number of situations, whereas ‘‘not(open(door))’’ is consistent with fewer: those in which the door is closed. However, such an extension of propositional theory begs the question of how meaning is computed in the first place. That is, how do we know which situations to which the proposition should be applied? And, if one considers that a descriptor such as ‘‘opened’’ behaves differently with different arguments (e.g., a bottle of soda is either opened or not opened, whereas a door may be wide open or slightly ajar), then the modified propositional theory approaches our embodied theory – that meaning depends on the implications for situated action. We revisit this topic in Section 3.2.3.

3.2. Experiment 2

Experiment 1 demonstrated that negated sentences are judged as situationally ambiguous compared to their positive counterparts. That is, the negated sentence corresponds to a greater variety of situations than does the positive sentence. Because the negated sentences are situationally ambiguous, they will not generally be useful for conveying new information (Givon, 1978). Nonetheless, negated sentences are used frequently. Givon (1978) suggests that in some pragmatic contexts negated sentences are no more ambiguous or difficult than the corresponding positive sentences. In particular, Givon proposes that felicitous use of negated sentences requires that they are used to address or counter presuppositions held by a listener/comprehender. We tested this prediction in Experiment 2.

Participants read positive and negated sentences similar to those used in Experiment 1, but this time the sentences were read in context. The contexts could be either supportive or non-supportive. Supportive contexts were designed to provide the pragmatic constraints Givon suggests are necessary for felicitous use of negated sentences. Thus, in the example given in Table 2, the supportive context mentions that choice of color of a new couch is important, and the critical sentence is either ‘‘The couch was black’’ or ‘‘The couch wasn’t black.’’ In the non-supportive context, the text again refers to a decision about the couch, but the dimension of choice (material) is different from that referred to in the critical sentence (color). To index comprehension

[1] For example, the topic of negation does not appear in the index of Van Dijk and Kintsch (1983) or in the index of Gernsbacher's *Handbook of Psycholinguistics* (Gernsbacher, 1994).

Table 2
Example of materials used in Experiment 2

Introduction	Marcy needed a new couch for her family room.
Supportive context	She wasn't sure if a darkly colored couch would look the best or a lighter color.
	She finally picked one out and had it delivered to her home.
Non-supportive context	She wasn't sure what kind of material she wanted the couch to be made of.
	She finally picked one out and had it delivered to her home.
Critical positive	The couch was black.
	It looked very nice in her family room.
Critical negative	The couch wasn't black.
	That probably would have been too dark.

difficulty, we measured the amount of time needed to read the critical sentence.

We can contrast the predictions of two accounts. Suppose, as propositional theory suggests, that negated sentences require extra processing (relative to positive sentences) to encode and comprehend. This extra processing is required to form the embedded propositional structure and to integrate the multiple propositions into the representation. This extra processing should slow down reading of negated sentences in both the supportive and non-supportive contexts. However, suppose Givon is correct. In this case, negated sentences should be felicitous in the supportive context and should be read about as quickly as the positive sentences. Only in the non-supportive contexts should the negated sentences be read more slowly than the positive sentences.

3.2.1. Method

Participants. Eighty students in an introductory psychology course participated for extra credit. All participants were native English speakers.

Materials and design. Sixty sets of texts were created (see Table 2 for an example). Each set consisted of short paragraphs corresponding to the four conditions (supportive or non-supportive context crossed with positive or negated critical sentences). The first one or two sentences in a text introduced a situation. The next sentence conformed to the context manipulation. In the supportive context, a dimension (e.g., color) was mentioned. This same dimension would be referred to in the critical sentences. In the non-supportive context, a different dimension (e.g., material) was mentioned. The context sentences were followed by a positive or negated critical sentence. Each critical sentence was followed by a closing sentence (e.g., ''That probably would have been too dark.''). The closing sentence was an attempt to reduce the bizarre nature of the negative sentence following the non-supportive context.

Eight of the sixty sets were modeled on the eight end-dimension sentences from Experiment 1. For example, in Experiment 1 a positive end-dimension sentence was ''The buttons are black,'' whereas in Experiment 2 a positive critical sentence was ''The couch was black.'' The mid-dimension terms from Experiment 1 were used in the supportive context condition to introduce the dimension later referred to by the critical sentence. The other sets of sentences were then modeled on these eight, although there was quite a bit of variation.

Each participant read sixty texts presented so that there were fifteen in each condition. Choice of condition for a particular text and order of presentation was counterbalanced over the participants for a total of eight counterbalance conditions. The participant read the text one sentence at a time (by hitting a key on the computer keyboard to advance through the text), and the times to read the sentences were recorded. Following each text, the participant was directed to type a few words describing the main topic of the text. This task was included to encourage careful reading.

3.2.2. Results

Because negated sentences always had one more syllable than their positive counterparts, the following procedure was used to analyze the data. First, slope and intercept coefficients for a linear regression of sentence reading times on the number of syllables were estimated for each condition separately for each participant. The regression was performed after centering the predictors (number of syllables) around zero by subtracting the average number of syllables from each predictor. In this way, the slope and intercept estimates were orthogonal, and the intercept estimates are the expected reading times for a sentence of average length in the experiment (average length was 7.50 syllables). Next, these estimates were analyzed as a mixed-design ANOVA with negation and context as within-subject factors and counterbalance condition as a between-subject factor. The counterbalance factor is included to optimize power and will not be discussed further (Pollatsek & Well, 1995).

The analysis of the average reading times (the intercepts) revealed a significant interaction between the two factors, $F1(1, 72) = 4.30$ ($F1$ is an analysis using differences among participants to compute the error term). In the unsupported condition, the average reading time of the negated sentences (1827 msec) was significantly longer than for the positive sentences (1673 msec), $F1(1, 72) = 15.42$. In the supported condition, however, the difference was eliminated: The average reading time for the negated sentences was 1548 msec, whereas the average for the positive sentences was 1607 msec, $F1(1, 72) = 2.28$, $p > 0.12$. Predictably, participants took much longer to read the critical sentences in unsupported texts (1750 msec) than in supported texts (1577 msec), $F1(1, 72) = 62.75$.

For the corresponding analysis using differences among items as a random factor, the analysis is slightly different because sentence length varies between items whereas participants read sentences of varying lengths within each treatment condition. First, we performed a regression analysis on mean reading times for items using number of syllables as the covariate. We then examined the residuals because these residuals have effects of sentence length statistically removed. This is equivalent to an ANCOVA, except that it allows for analysis of within-items manipulations (negation and context). The ANOVA on the residuals revealed a significant interaction between negation and context, $F2(1, 59) = 6.47$ ($F2$ uses item variability to compute the error term). In the unsupported condition, the effect of negation was quite strong, $F2(1, 59) = 25.23$, whereas in the supported condition the effect of negation was not significant $F2(1, 59) = 2.26$, $p = 0.138$.

No significant differences were detected for the effects of negation for the sentences preceding or following the critical sentences.

3.2.3. Discussion

Contrary to the predictions of propositional theory, readers processed negated sentences as easily as positive sentences when the pragmatic constraints for negation were satisfied. Only when texts violated the pragmatic constraints for negation were negated sentences more difficult to process. Thus, contrary to propositional theory, negated sentences are not always more difficult to comprehend than their positive counterparts. Why then has so much previous research (e.g., Carpenter & Just, 1975) led to the conclusion that negatives are difficult and require extra processing? Because the previous research presented sentences without regard to pragmatic constraints, and so the negated sentences were almost certainly viewed as situationally ambiguous (as in Experiment 1), and thus difficult to understand.

3.3. Experiment 3

The contrast between the results of Experiment 2 (negated sentences need not be difficult to process) and previous results using the sentence–picture verification procedure (negated sentences are difficult to process) is consistent with two explanations. First, as we suggested above, difficulty with negated sentences may arise only when the procedure violates pragmatic constraints on the use of negated sentences. Second, the materials that we used in Experiments 1 and 2 may be odd in such a way as to eliminate the difference between negated and positive sentences. To decide between these accounts, we used sentences from Experiments 1 and 2 in a sentence–picture verification task. Thus, if it is the violation of pragmatic constraints that is important, we should see a large difference between negated and positive sentences in the sentence–picture verification task even when using the sentences from Experiments 1 and 2. On the other hand, if the sentences used in Experiments 1 and 2 are odd, then we should see little or no difference between negated and positive sentences in the sentence–picture verification task.

3.3.1. Method

Participants. Twenty-eight native English speakers recruited from the campus community were paid five dollars to participate.

Materials and design. We used the eight sets of sentences constructed for Experiment 1. Two pictures were drawn for each set. One picture was consistent with the positive sentences and one was consistent with the negated sentences. Each block consisted of sixty-four trials corresponding to the eight sets by four sentences in a set by two pictures. Note that pairing a sentence with a consistent picture created a ‘‘true’’ trial and that pairing a sentence the alternative picture created a ‘‘false’’ trial. The order of trials was randomized for each participant so that within each sub-block of eight trials there was one sentence from each set and four false and four true trials.

Procedure. Participants responded to five blocks of sixty-four trials. The first block was considered practice. On each trial the selected sentence was displayed on a computer monitor and the participant pressed the space bar on the computer keyboard when finished reading the sentence. The reading time between presentation of the sentence and the key press was recorded. Next, the appropriate picture was presented and the participant indicated whether the sentence described the picture by pressing one of two keys. The time elapsed between presentation of the picture and the response was the decision time. Half of the participants pressed the ‘‘true’’ key with their dominant hand and half pressed the ‘‘false’’ key with their dominant hand. A low tone was sounded if the participant’s response was incorrect.

3.3.2. Results and discussion

For both the reading times and the decision times the data were analyzed as follows. First, the data from the first block were treated as practice, leaving four additional blocks for analysis. Second, reaction times associated with incorrect responses were eliminated. Third, for each participant and each condition, reaction times that were more than 2.5 standard deviations from the mean for that participant and condition were eliminated. For the remaining data, means for the conditions are presented in Table 3. Analyses were also conducted using sentence set as a random factor and reported as $F2$. For these analyses, reaction times were trimmed by 2.5 standard deviation from the mean for that sentence set and presentation condition.

Table 3
Reading times and decision times (in ms) from Experiment 3

	End-dimension		Mid-dimension	
	Positive	Negated	Positive	Negated
Reading times				
True	657	807	684	823
False	652	807	680	852
Decision times				
True	906	1180	902	1221
False	1016	1344	1030	1310

Reading times. If our sentences are comparable to those used in previous work on negation, then the negated sentences should take longer to read than the positive sentences. Indeed, this was the case. The difference in reading time between negated sentences (822 msec) and positive sentences (668 msec) was significant, $F1(1, 27) = 16.41$, $F2(1, 7) = 280.20$. The mid-dimension sentences (760 msec) took longer to read than the end-dimension sentences (731 msec), $F1(1, 27) = 6.06$ (but $F2(1, 7) = 1.04$), probably because the mid-dimension sentences are less situationally specific than the end dimension sentences (see results of Experiment 1). There was little difference in time to read the true (743 msec) and false (748 msec) sentences, $F1 < 1.0$, $F2 < 1.0$), because the variable was not relevant until after the picture was presented.

Decision times. Decisions were much slower for negated sentences (1264 msec) than for the positive sentences (964 msec), $F1(1, 27) = 128.25$, $F2(1, 7) = 171.16$. Also, decisions were slower for false sentences (1175 msec) than for true sentences (1052 msec), $F1(1, 27) = 74.29$, $F2(1, 7) = 71.62$. The effect of dimension was not significant, $F1 < 1$, $F2 < 1$. There was, however, a marginally significant three factor interaction, $F1(1, 27) = 3.75$, $p = 0.063$, $F2(1, 7) = 3.32$, $p = 0.11$. For the end-dimension, the difference between true and false is a bit smaller for positive sentences (110 msec) than for negated sentences (164 msec), whereas for the mid-dimension, the difference between true and false is a bit larger for the positive sentences (128 msec) than for the negated sentences (89 msec).

On average, participants were correct on 94% of the trials. There was, however, a significant difference in error rate between the positive sentences (97%) and the negated sentences (91%), $F1(1, 27) = 58.72$.

The conclusion is clear. When our sentences are presented in isolation in the sentence–picture verification procedure, there appears to be a large cost to negation. This cost is in stark contrast to the absence of any cost when the sentences are presented in a supporting context, as in Experiment 2.

3.4. General discussion

We have demonstrated two important facts. First, out of context, negated sentences are often interpreted as situationally more ambiguous than the corresponding affirmatives. Second, when the negated sentences are put into an appropriate context, they are no more difficult to process than the corresponding affirmatives. Both of these facts speak against the theory that people understand language in terms of propositions. According to published propositional theories, a negated sentence is understood as a negated proposition. The information content of negated propositions is equivalent to (but reversed from) that of the corresponding positive proposition. Thus, these theories predict that negated sentences should be understood as specific as the corresponding affirmatives, and this prediction is contradicted by our first fact. Because negated propositions are more complex than the corresponding positive propositions, propositional theory predicts that negated sentences should be difficult to understand and to integrate with a text, and this prediction is contradicted by our second fact.

We believe these two facts are related. Negated sentences are generally not useful for conveying new information because of the situational ambiguity. That is, when attempting to use a negated sentence to convey new information, the information is ambiguous (the first fact) and hence is difficult to process. Nonetheless, negated sentences are useful for updating or clarifying presupposed or given information, and hence are processed felicitously in context (the second fact). In this way, the second fact can be seen as a consequence of the first fact: Negated sentences require appropriate contexts because they are ambiguous out of context. Other common uses of negation in discourse, such as irony and politeness, may also trade on the situational ambiguity associated with negatives (Giora, 1995).

These data speak against the propositional account of negated sentences, but is that sufficient evidence to discount all of propositional theory? Clearly not, but there other reasons to be suspicious of propositional theory. Barsalou (1993), Glenberg (1997), Glenberg and Robertson (in press (b)), Lakoff (1987), and Shanon (1988) lay out these reasons. Here is a brief overview.

The other empirical bases for propositional theory (recall of gist and priming relations) do not provide strong support. Memory for text is often gist-like, rather than verbatim (Kintsch, 1974; Sachs, 1967). This finding can be accommodated by any theory that eschews a verbatim representation, however. For example, suppose that memory for text consists of images, perceptual symbols (Barsalou, 1993), or a meshed set of affordances (Glenberg & Robertson, in press (a), (b)) that do not have a one-to-one correspondence with sentences. Then verbal recall will be gist-like as these representations are translated into language for the purpose of recall. McKoon and Ratcliff (1980) demonstrated that concepts prime one another more strongly when they appear within propositional units (or in separate propositions that are themselves closely related). This finding can also be accommodated by any theory that proposes that memory for text is structured along dimensions of meaning. That structure, however, need not be propositional.

Propositions such as ‘‘next-to (printer, monitor)’’ and ‘‘on (monitor, table)’’ seem to be direct representations of the corresponding ideas, but that is only because many problems have been glossed over. The first problem is that in a theory of the meaning of language, the arguments cannot be words (e.g., ‘‘monitor’’), because the words themselves require definitions. One solution is to propose that the arguments are not really words (although we may use words for notational convenience). Instead, the arguments are abstract, amodal, arbitrary symbols, that point to the real meaning of the words in a lexicon or semantic space. But, what is this real meaning? Some have proposed that meaning of a concept is a list of features. Barsalou (1993) and Shanon (1988) present compelling arguments against this idea. Barsalou demonstrates that a list of features does not capture important relations. For example, the definition of ''chair'' cannot be the list of features ''back'' ''seat'' and ''four legs,'' because a back, a seat, and four legs in a pile on the floor does not make a chair: the components have to be related appropriately. Those relations may be represented by additional propositions (remember, propositions are designed to relate), but that doesn't solve the initial problem of how we get meaning for the arguments being related.

A second solution proposed for the meaning of arguments is a semantic network. Thus, the meaning of ''monitor'' comes about from all of its relations to other words in a semantic network. In such a network a monitor is the thing that is (a) part of a computer system, (b) has a display screen, (c) is like a TV set, etc. Keep in mind, however, that all of these terms require definition by their relations to other terms. Can everything be defined in terms of something else? The answer is almost certainly not. Imagine, for example, that you are stranded in a country whose residents use a language you do not know. At your disposal is a dictionary written totally in the unknown language. When reading a sign, you could use the dictionary to look up the first word of the sign and see its definition. But that definition requires that you look up more words in the language, and each of those words requires that you look up still more. Clearly, the meaning of the sign would never become evident using this method. This is the symbol grounding problem (Harnad, 1990; Lakoff, 1987; Searle, 1980).

A third solution is to propose that the arguments take on meaning by referring to objects and concepts outside of the system. That is, we need perception to assign meaning to internal abstract, amodal, arbitrary symbols. While easy to say, this account has tremendous difficulties. Lakoff has documented some of these difficulties. Among them are (a) many concepts do not correspond to simple elements or combination of elements in the physical world, (b) many elements and combinations of elements in the physical world do not map onto concepts, and (c) it may be impossible in principle to find the one correct mapping from the system of elements in the world to the system of symbols in a cognitive system. At this point the reader may object by noting that we *do* know the meaning of linguistic expressions and clearly we can accurately name many objects in the world, so something must be wrong with Lakoff's analysis. Keep in mind, however, that Lakoff's analysis does not demonstrate the impossibility of meaning. It demonstrates the impossibility of finding the correct meaning for the abstract, amodal symbols arbitrarily related to their referents presumed to underlie propositions. Thus, the problem may not be with Lakoff or us, but with propositions.

The problems for propositions do not end with the problems for abstract, amodal, arbitrary symbols used for arguments. A second problem is in the workings of the relational terms. The meaning of ''on (monitor, table)'' only seems self-evident because we, the readers, supply the meaning upon reading the word ''on.'' But in a theory of meaning, how is the meaning of the relational term captured? The only solution offered within propositional theory is to have extrinsic rules of inference or interpretation (Johnson-Laird, 1983; Kintsch, 1974). That is, a set of rules or productions (e.g., if next-to (X, Y), then next-to (Y, X)) specify the interpretation and use of relations.

Once again, there are serious flaws in this system. First, defining relations in terms of other relations does not seem to be any better than defining abstract symbols in terms of other abstract symbols. Second, the meaning of common words such as ''on'' often depends on the situation, not the words. For example, the meaning of ''on'' in ''The monitor is on the table'' (here ''on'' means sitting on top of a horizontal surface) is quite different from the meaning of ''on'' in ''The monitor is on the wall'' (here ''on'' means attached to a vertical surface). Of course, one could assert that there are two relations, ''on-1'' and ''on-2.'' But, forming the right propositions using ''on-1'' or ''on-2'' requires understanding the situation, and if the situation is understood *before* the proposition is formed, what is the role of the proposition (Barton & Sanford, 1993)?

Some theorists have attempted to get around these problems by building into their theories ''world knowledge'' that is called upon to resolve the problems of knowing which propositions to build. For example, Kintsch's construction–integration model (Kintsch, 1988) works in part by forming many propositions (the construction phase), some of which are retrieved from memory (world knowledge). There will be incompatibilities between some of the propositions (e.g., the two meanings of ‘‘on’’), and incompatible propositions are connected by inhibitory links. The propositions are trimmed by applying a constraint satisfaction algorithm that allows a set of mutually consistent propositions to be foregrounded. There are three problems with this proposed solution. First, artificial systems based on vast compilations of propositional facts do not seem to work. That is, they become extremely unwieldy as the number of facts begin to accumulate (Winograd & Flores, 1987). Second, adding more and more propositions formed from abstract, amodal, arbitrary symbols defined solely in terms of other abstract symbols does not the solve the original problem: How do any of those propositions become meaningful?

Finally, abstract, amodal, arbitrary symbols are insufficient *in principle* to account for our ability to understand language about novel situations. Because the symbols are arbitrarily related to the objects, there can be no emergent properties. That is, all knowledge of the objects must be directly represented in the propositions or logically (computationally) derivable from the propositions. That this is false is easy to demonstrate when considering language about novel situations (Glenberg & Robertson, in press (b)). Consider the two versions of the following sentence: ‘‘After wading barefoot in the lake, Eric dried his feet using his shirt/glasses.’’ From the point of view of abstract, amodal, arbitrary symbols, the versions of the sentence are equally sensible: Both versions are grammatically correct, both respect semantic selectional constraints such as animacy, and both can be easily propositionalized. Also, considering world knowledge, ‘‘dried’’ is equally unrelated to ‘‘shirt’’ and ‘‘glasses’’ (see Glenberg & Robertson, in press (b), for support of this claim). But clearly the sentences are not equally sensible! Consider another example of the same point: ‘‘Marissa forgot to bring her pillow on a camping trip; as a substitute, she filled up an old sweater with leaves/water.’’ Again both sentences are grammatical, both can be propositionalized, and considering world knowledge, sweaters are more strongly (and with positive valence) associated with water than with leaves. Nonetheless, the sentence with ‘‘water’’ is nonsense, because a water-filled sweater does not make a pillow. Theories of meaning based on abstract, amodal, arbitrary symbols cannot account for the differential sensibility of these sentences because the theories have no access to what sweaters or leaves really are; they have access only to the abstract symbols. Hence if it is not specifically encoded that (the symbol for) sweaters can be stuffed with (the symbol for) leaves to make (the symbol for) pillows, the theory cannot derive it because the theory has no access to real sweaters and how they might interact with other real substances.[2] In short, theories of meaning based on abstract, amodal, arbitrary symbols cannot discriminate between sensible and non-sensible descriptions of novel situations.

So, how do we understand what sentences mean? The key seems to be to make use of symbols that are not abstract, amodal, and arbitrarily related to their referents, and hence to drop the idea of a proposition, or at least propositions as we know them. Examples of the use of non-arbitrary symbols are Lakoff’s and Mandler’s image schemas and Barsalou’s perceptual symbols (Lakoff 1987, Mandler 1992, Barsalou 1993). Here we focus on an account developed by Glenberg (1997); Glenberg and Robertson, in press (a), (b)).

Glenberg (1997) claims that cognitive systems co-evolve with bodies to facilitate effective action, that is, action appropriate for specific types of bodies. For example, if a human tried to escape a predator by flapping his or her arms to fly away (so effective for a bird’s body), it is unlikely that human would survive long enough to contribute to the gene pool. Thus, cognition must contribute to effective action if cognition is to evolve. Just as clearly, cognition is the source of meaning: how we construe situations. On this analysis, the cognitive system

[2] One might argue that the proposition ‘‘Wet things do not make good pillows’’ is logically derivable from other knowledge. There are at least three problems with this claim, however. First, a long enough chain of fuzzy inferences could also derive that leaves (which are often wet, dirty, and decayed) do not make good pillows. Second, people do not appear to engage in such inferential reasoning when understanding these sentences. The comprehension time for sentences such as ‘‘Eric dried his feet using his shirt’’ is no different from the time to comprehend ‘‘Eric dried his feet using his towel.’’ Third, some wet things, such as the pillow on a floating mat in a lake, make really fine pillows.

provides a link between embodied action and meaning. That link is captured within an action-based definition of meaning (Glenberg, 1997): ''The meaning of a situation corresponds to the set of actions available to an individual in that situation. This set of actions results from the mesh of affordances, knowledge, and goals.''

Affordances (cf. Gibson, 1979) are possibilities for interaction jointly determined by bodily capabilities and physical aspects of the environment. For example, for an adult, a chair affords sitting, but it also affords standing-on to change a light bulb, and by lifting the chair into a defensive position, it affords protection against snarling dogs. The same chair might afford sitting for a toddler, but because the toddler cannot lift the chair into a defensive position, the chair does not afford defense for the toddler. A second component of the set of actions is action-based knowledge encoded in memory. To continue the example, our knowledge that an old chair is rickety changes the meaning of the chair so that it no longer affords standing-on. These various determinants of action are meshed, that is, integrated in a way that respects the constraints of physics and biology. Thus, we can mesh the affordances of a chair with the goal of raising the body to change a light bulb; but it is difficult or impossible to mesh the affordances of a chair with the goals of, say, swimming across a lake, or jumping rope. Human bodies simply do not work in a way that meshes these affordances and these goals. Thus, the meaning of a situation involving an adult, a kitchen, a chair, and a blown light bulb arises from the mesh of affordances (for the adult) of the chair in the kitchen, the individual's action-based knowledge of the chair, and the goal of changing the bulb. The meaning of the situation would be different for a toddler or for a dog.

This embodied account of meaning is related to language through the Indexical Hypothesis (Glenberg & Robertson, in press (a), (b)). According to this hypothesis, language comprehension results from three steps. The first is to index words and phrases to objects or analogical representations (e.g., perceptual symbols; Barsalou, 1993) of objects. The second step is to derive affordances from the objects. The third step is to mesh the affordances under the guidance of syntax (Kaschak & Glenberg, 1999). Consider again the sentences ''After wading barefoot in the lake, Eric dried his feet using his shirt/glasses.'' The phrases ''his shirt'' and ''his feet'' are indexed to actual objects or their perceptual symbols; the affordances of shirts and feet are derived; under the guidance of syntax (e.g., Eric is the actor, his feet are the object, his shirt is the instrument, and the transitive form of the sentence specifies that the actor acts on the object) the affordances are meshed to accomplish the goal of drying. Similarly, the affordances of sweaters and leaves can be meshed to accomplish the goal of making a pillow for Marissa. We do not need to have previously learned that sweaters can be stuffed with leaves to make pillows. Instead, the sensibility of the actions (and hence the sentence) emerges from the mesh of the affordances guided by syntax.

Words and syntax alone do not produce meaning. In fact, the importance of indexing words to objects (and deriving affordances) is easily demonstrated. First, when affordances do not mesh (the affordances of glasses and feet cannot be meshed to accomplish the goal of drying), then sentences such as ''Eric dried his feet using his glasses'' do not make sense. But second, a change in referent, and hence a change in the affordances, can make the same words and syntax perfectly sensible. Thus, if Eric had been clowning around and was wearing a pair of large spongy glasses, then the sentence ''Eric dried his feet using his glasses'' is easy to understand. The point is that it is not just the properties of words or abstract symbols that determine the meaningfulness of sentences, instead, meaning is determined by the affordances of the objects to which the words refer.[3]

This embodied account of cognition helps us to

[3]From these examples, it may appear that an embodied account of language can only address language about concrete entities. Although we have not worked out a compelling account of abstract ideas within the Indexical Hypothesis, other embodied approaches to cognition have. Lakoff (1987) suggests that abstract ideas (e.g., ''argument'') are understood by metaphor using concrete ideas (e.g., an argument is like a building). That, he suggests, is why our language about abstract ideas is often concrete and systematically related to the metaphor (e.g., ''The argument has a complex structure,'' ''The argument collapsed,'' ''The argument is supported by the data,'' and so on). Barsalou (in press) demonstrates how many abstract ideas (e.g., ''truth,'' ''exclusive or'') can arise from the comparison of perceptual symbols to the observed world.

understand why negated sentences appear ambiguous or uninformative when presented out of context. Consider the prototypical sentence from a sentence–picture verification experiment, ‘‘The star is above the cross,’’ or one with a bit more content, ‘‘The cat is in the tree.’’ ‘‘The cat’’ can be indexed to an actual cat or a perceptual symbol of a cat, and affordances derived. Similarly ‘‘the tree’’ can be indexed and affordances derived. These affordances can be meshed as guided by syntax to represent the cat in the tree (rather than, say, vice versa). The meshed conceptualization facilitates action in that it tells us how to act if the goal is to engage the cat (go to the tree), and it tells us how to act if the goal is to avoid the cat (avoid the tree). Now consider the negated sentence, ‘‘The cat is not in the tree.’’ Although we can derive affordances from the perceptual symbols for a cat and a tree, the syntax of the sentence gives no instruction in how to mesh the two; in fact the syntax directs us to avoid meshing the two. In this case, we have no coherent, that is, meshed, representation of the situation that can be used to guide action; instead, there is a cat, and separately there is a tree. Thus, out of context, the negated sentence is ambiguous or uninformative in the sense that it provides little information that can be used to guide action. In a supportive context, however, the positive and the negated sentences can be equally informative, because both result in a meshed conceptualization. For example, in the context ‘‘Either the cat is in the tree or the cat is in the den,’’ the positive sentence ‘‘The cat is in the tree’’ instructs the comprehender to mesh affordances of cat and tree. The negated sentence, ‘‘The cat is not in the tree (but in the den),’’ produces a mesh of cat and den.

4. Conclusions

Our experimental results have removed one of the pillars of support for propositional theory: slow comprehension of negated sentences. That alone is not enough to topple the theory, but it is enough to motivate serious thought about replacing propositions as the basic elements of thought. Those who are committed to the theory of propositions built from abstract, amodal, arbitrary symbols will not be greatly discomforted by these findings. A proponent of propositions might argue that (a) propositions work well enough, (b) cognitive scientists have already introduced correctives such as images and mental models, (c) everyone knows that propositions are a convenient fiction, a shorthand, for whatever is doing the real work of representation and meaning, and (d) there are no good alternatives anyway. But, none of these arguments are valid. (a) Propositions only work well enough in laboratory tasks where there is no need to model meaning, only syntactic manipulation. (b) Many correctives to simple propositions are themselves constructed from propositions. For example, several versions of mental model theory rely on situation models constructed from propositions (Kintsch, 1988; Van den Broek, 1994). (c) By treating propositions as a convenient fiction we abrogate our responsibility as scientists. Instead of searching for causes we become satisfied with stories. Also, the commitment to propositions keeps us looking at a very limited range of phenomena and tasks: only those that fit the mold of propositional analysis. When we examine language about novel situations (e.g., Marissa and her pillow), the vacuousness of propositions composed of arbitrary symbols is quickly revealed. (d) There are alternatives, although admittedly they are not yet as well developed as propositional theory.

We have sketched one alternative based on the concept of embodiment. Embodied accounts are appealing because they hold the promise of coming to grips with meaning and how meaning contributes to action. But, those advances come at the cost of dropping two assumptions about cognition: that it is based on abstract, amodal, arbitrary symbols, and that those symbols are used in a process of formal computation. As the history of science has demonstrated, there is much to be gained by formalizing theories, and that is likely to be true in cognitive science as well. But, as we have demonstrated here, there is much to be lost by presuming that the process of cognition is itself constituted of formal operations on abstract symbols.

Acknowledgements

A University of Wisconsin Graduate School Research Committee grant (990288) to Arthur Glenberg and a National Institute of Mental Health (T32-

MH18931) predoctoral traineeship to David Robertson supported preparation of this report.

References

Barsalou, L. W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In: Collins, A. C., Gathercole, S. E., & Conway, M. A. (Eds.), Theories of memories, Erlbaum, London.

Barsalou, L. W., Perceptual symbol systems. *Behavioral and Brain Sciences*, in press.

Barton, S. B., & Sanford, A. J. (1993). A case study of anomaly detection: Shallow semantic processing and cohesion establishment. *Memory & Cognition 21*, 477–487.

Carpenter, P. A., & Just, M. A. (1975). Sentence comprehension: A psycholinguistic processing model of verification. *Psychological Review 82*, 45–73.

Carpenter, P. A., & Just, M. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review 99*, 122–149.

Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior 12*, 335–359.

Clark, H. H., & Chase, W. G. (1972). On the process of comparing sentences against pictures. *Cognitive Psychology 3*, 472–517.

De Villiers, J. G., & Flusberg, H. B. T. (1975). Some facts one simply cannot deny. *Journal of Child Language 2*, 279–286.

Gernsbacher, M. A. (1994). Handbook of Psycholinguistics, Academic Press, San Diego.

Gibson, J. J. (1979). The ecological approach to visual perception, Houghton Mifflin, New York.

Giora, R. (1995). On irony and negation. *Discourse Processes 19*, 239–264.

Givon, T. (1978). Negation in language: Pragmatics, function, ontology. In: Cole, P. (Ed.), Syntax and semantics 9: Pragmatics, Academic Press, New York.

Glenberg, A. M. (1997). What memory is for. *Behavioral and Brain Sciences 20*, 1–55.

Glenberg, A. M., Robertson, D. A., Indexical understanding of instructions. *Discourse Processes*, in press.

Glenberg, A. M., Robertson, D. A., Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language*, in press.

Harnad, S. (1990). The symbol grounding problem. *Physica D 42*, 335–346.

Horn, L. R. (1989). A natural history of negation, University of Chicago Press, Chicago.

Johnson-Laird, P. N. (1983). Mental models, Cambridge, MA. Harvard University Press.

Kaschak, M. P. & Glenberg, A. M. (1999). Constructing meaning: The role of affordances and grammatical constructions in sentence comprehension. Submitted for publication.

Kintsch, W. (1974). The representation of meaning, Erlbaum, Hillsdale, NJ.

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

Lakoff, G. (1987). Women, fire, and dangerous things: What categories reveal about the mind, University of Chicago Press, Chicago.

MacDonald, M. C., & Just, M. A. (1989). Changes in activation level with negation. *Journal of Experimental Psychology: Learning, Memory and Cognition 15*(4), 633–642.

Mandler, J. M. (1992). How to build a baby: II. Conceptual primitives. *Psychological Review. 99*, 587–604.

McKoon, G., & Ratcliff, R. (1980). Priming in item recognition: The organization of propositions in memory for text. *Journal of Verbal Learning and Verbal Behavior 19*, 369–386.

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.

Pylyshyn, Z. W. (1973). What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin 80*, 1–24.

Robertson, D. A., Glenberg, A. M., & Members of the Honors Seminar in Cognitive Psychology (1999), Grounding symbols and computing meaning: A supplement to Glenberg & Robertson. Manuscript. http://psych.wisc.edu/glenberg/supplement_jml_g&r.html.

Sachs, J. D. S. (1967). Recognition memory for syntactic and semantic aspects of connected discourse. *Perception & Psychophysics 2*, 437–442.

Satterthwaite, F. W. (1946). An approximate distribution of estimates of variance components. *Biometrics Bulletin 2*, 110–114.

Searle, J. R. (1980). Minds, brains and programs. *Behavioral and Brain Sciences 3*, 417–424.

Shanon, B. (1988). Semantic representation of meaning: A critique. *Psychological Bulletin 104*, 70–83.

Tanenhaus, M. K., Carroll, J. M., & Bever, T. G. (1976). Sentence–Picture verification models as theories of sentence comprehension: A critique of Carpenter and Just. *Psychological Review 83*, 310–317.

Trabasso, T., Rollins, H., & Shaughnessy, E. (1971). Storage and verification stages in processing concepts. *Cognitive Psychology 2*, 239–289.

Van den Broek, P. (1994). Comprehension and memory of narrative texts. In: Gernsbacher, M. A. (Ed.), Handbook of psycholinguistics, Academic Press, New York, pp. 539–588.

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

Wason, P. C. (1965). The contexts of plausible denial. *Journal of Verbal Learning and Verbal Behavior 4*, 7–11.

Winograd, T., & Flores, F. (1987). Understanding computers and cognition: A new foundation for design, Addison-Wesley, Reading, MA.