METAPHOR AND SYMBOL, *16*(3&4), 173–192

Salience and Context Effects: Two Are Better Than One

Orna Peleg and Rachel Giora
Department of Linguistics
Tel Aviv University

Ofer Fein
Department of Behavioral Sciences
The Academic College of Tel-Aviv-Yaffo

This study provides evidence supporting the hypothesis that language comprehension involves 2 separate mechanisms that run in parallel: a linguistic mechanism and a contextual mechanism. The linguistic mechanism (e.g., lexical access) is modular and stimulus driven; it is a bottom-up, perceptual mechanism, induced by a lexical stimulus to search the mental lexicon for its match. This mechanism is encapsulated with respect to nonlinguistic information and thus operates locally (i.e., on the word level). Lexical access is exhaustive and ordered: Salient meanings are accessed faster.

Contextual facilitation, on the other hand, is the outcome of a central, expectation-driven mechanism that operates globally during language comprehension at the point where prior linguistic information has already been processed and interfaced with other cognitive processes (e.g., inferencing).

Experiment 1 indicates that contextual facilitation can occur even before lexical accessing takes place, fostering an impression of a selective process. Experiment 2 shows that the target word's position in the sentence (initial vs. noninitial) is crucial for the operation of the predictive mechanism. Thus, we would not expect contextual meanings to outweigh salient meanings at the beginning of sentences.

For more than 2 decades, linguists and cognitive scientists have looked into how and when contextual information affects initial processing. In this article, we also wish to shed light, primarily empirically, on how salient meanings and senses of

Request for reprints should be sent to Rachel Giora, Department of Linguistics, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: giorar@post.tau.ac.il

words shape our psycholinguistic behavior vis à vis contextual information. Particularly, we focus on "false positives" (Fodor, 1983)—meanings that are contextually incompatible but are activated nonetheless.

Two major approaches to initial processes have dominated the field since the 1970s—the direct access model view and the modular view.

THE DIRECT ACCESS VIEW

The direct access or interactionist view, held by linguists (e.g., Carston, 1999; Sperber & Wilson, 1986/1995), cognitive scientists (e.g., Martin, Vu, Kellas, & Metcalf, 1999; Vu, Kellas, Metcalf, & Herman, 2000; Vu, Kellas, & Paul, 1998), and psycholinguists (see Gibbs, 1994, and Giora, in press-b for a review), attributes to context a crucial role in initial processes. This view assumes a single interactive mechanism, which is sensitive to both linguistic and nonlinguistic information. In this view, contextual information interacts with the lexicon very early on and results in selective access of contextually appropriate meanings. Particularly, in a rich and strongly supportive context, the contextually appropriate meaning is tapped initially, directly, and exclusively, without involving any contextually incompatible phase at all. Rather, contextually incompatible meanings are inhibited and hence fail to reach sufficient levels of activation (cf. Martin et al., 1999; Vu et al., 1998).

THE MODULAR VIEW

The modular view (e.g., Fodor, 1983) posits a mechanism sensitive only to linguistic information. Lexical accessing is modular: stimulus driven, invariant across contexts, encapsulated (impenetrable to processes occurring outside the input system), automatic and fast, and, on some traditional views, also exhaustive and unordered: *All* the meanings of a word are activated on its encounter, regardless of either contextual bias or frequency (e.g., Connine, Blasko, & Wang, 1994; Onifer & Swinney, 1981; Swinney, 1979; Till, Mross, & Kintsch, 1988; and see Giora, in press-b, for a review). On this view, initial processes would (also) involve contextually inappropriate meanings that will have to be discarded postlexically. Contextual processes, then, affect comprehension only *after* all the meanings have been activated.

Findings, however, have not been monolithic. On the one hand, comprehension was shown to involve contextually incompatible meanings even when context was strongly biased in favor of one sense (e.g., Gibbs, 1980, 1990; Giora, 1997, 1999; Rayner, Pacht, & Duffy, 1994; Swinney, 1979). Such findings are consistent with the modular view but disconfirm the direct access view. On the other hand, recent studies by Vu, Kellas, and their colleagues (Martin et al., 1999; Vu et al., 1998, 2000) have shown that a strong context affects comprehension directly. Contexts bi-

ased toward any one meaning of an ambiguous word activated that meaning exclusively. Such findings support the direct access view but question the modular view.

THE REORDERED ACCESS MODEL

To account for the conflicting findings, Rayner and his colleagues (e.g., Binder & Rayner, 1998, 1999; Duffy, Morris, & Rayner, 1988; Rayner et al., 1994; and see also Kawamoto, 1993) proposed a hybrid model, according to which context can interact with the lexicon and affect lexical access by boosting the activation levels of the contextually compatible meaning. Diverging from the direct access view, however, these enhancing effects are not inhibitory: Although the contextually compatible meaning is facilitated, the levels of activation of the alternative meaning(s) are not reduced. Thus, when context is biased in favor of a dominant meaning of an ambiguous word (e.g., the institutional meaning of *bank*) or in favor of one of the meanings of a balanced ambiguity (where some meanings are equally salient, e.g., *bug*), the levels of activation of that meaning are elevated, resulting in its fast and direct access. Similarly, when context is biased in favor of the less salient, subordinate meaning of an ambiguous word (e.g., the riverside meaning of *bank*), that meaning reaches sufficient levels of activation and is consequently processed alongside the incompatible dominant meanings whose high accessibility is still intact. The interference of the incompatible dominant meaning, however, slows down lexical processes. Thus, whereas the reordered access view pairs with the direct access model in predicting early interactive context effects, it differs from the direct access view on the inhibition assumption. Both, however, assume a single mechanism that is sensitive to both linguistic and nonlinguistic information.

THE GRADED SALIENCE HYPOTHESIS

Another attempt to account for the conflicting findings has been made by the graded salience hypothesis (Giora, 1997, 1999, in press-b; Giora & Fein, 1999; Giora, Fein, & Schwartz, 1998; Peleg, 2000). The graded salience hypothesis assumes two different mechanisms, one sensitive to linguistic information and one sensitive to linguistic and nonlinguistic information (cf. Giora, in press-b; Peleg, 2000). Accordingly, like the modular view (but diverging from the alternative models), the mechanism responsible for lexical access is sensitive only to information included in the mental lexicon. However, unlike the traditional modular assumption (Fodor, 1983), the graded salience hypothesis further assumes that lexical access is salience oriented: Salient meanings are activated before less salient ones (see also the ordered access view, e.g., Duffy et al., 1988; Rayner & Frazier, 1989; Rayner & Morris, 1991; Sereno, Pacht, & Rayner, 1992; and see Giora, in press-b, for a review).

To be salient, the meanings of a word or an expression have to be stored or coded in the mental lexicon and be foremost on our mind due to conventionality, frequency, familiarity, or prototypicality. Meanings not coded in the mental lexicon (e.g., conversational implicatures constructed on the fly) are nonsalient. Coded meanings that are less frequently used or are less familiar are less salient. Thus, for computer freaks, the computer sense of *mouse* must be more salient than the mammal sense, because they interact with the former more often than with the latter. A familiar irony (*"Read my lips"*) may have two similarly salient meanings—the literal and the ironic; an innovative use of it (*"Read my lipstick"*), however, may have a couple of salient meanings—the literal, the ironic, and a nonsalient meaning questioning the credibility of the speaker on account of her "femininity." Similarly, a familiar proverb (*"Lightning never strikes the same place twice"*) may have two salient meanings—the literal and the proverbial; an innovative proverb, however, has only one salient meaning—the literal (made up of the salient literal meanings of its components, cf. Katz & Ferretti, 2001/this issue).[1] Salient meanings are accessed via a direct look up in the mental lexicon, regardless of context. Thus, according to the graded salience hypothesis, initial processes may involve meanings activated on account of their salience rather than on the basis of their compatibility with context.

Recently, however the graded salience hypothesis has been challenged by studies by Kellas and his colleagues (Martin et al., 1999; Vu et al., 1998, 2000) claiming to have provided evidence in favor of selective access. We propose that the selective access demonstrated by these studies may be the result of a mechanism other than lexical accessing.

Following Fodor (1983, and see also Neely, 1977), the graded salience hypothesis assumes that contextual information may affect comprehension immediately. However, such processes do not interact with lexical accessing, but run in parallel. Our proposal is that contextual processes make up a distinct mechanism that has a predictive but not a blocking effect. Under certain conditions, this mechanism may avail the contextually appropriate concept immediately. It cannot, however, inhibit salient meanings activated independently by the lexical mechanism on encounter of the lexical stimulus. Indeed, contextual information may be strong and even have faster effects than lexical processes, so much so that it may avail appropriate interpretations even before the relevant stimulus is even encountered, fostering an impression of a selective access. This may be particularly true when the stimulus is placed at the end of a strong sentential context, allowing for guessing

[1]For a review of the literature as to when salient interpretations of phrases longer than a word are accessed see Giora (in press-b) and see also Frisson and Pickering (2001/this issue). On how the graded salience hypothesis accounts for a great range of linguistic phenomena such as ambiguity resolution, literal and nonliteral interpretations, and comprehension of nonsalient inferred interpretations, see Giora (in press-b).

and inferential processes to take place. However, it does not interact with lexical processes and selects the appropriate meaning, but operates simultaneously.

Space does not allow an in-depth discussion of the notion of predictiveness or strength of context. Suffice it to say, however, that a predictive context is one that is constraining enough to allow the guessing of an oncoming concept or meaning. Several factors may account for the predictiveness of contextual information. For instance, the strength of the association of cause and effect can index the predictability of the effect on the basis of the cause (see Keenan, 1978; Klin, Murray, Levy, & Guzmán, 1999). Similarly, manifesting features salient in the target concept (see, e.g., the contexts in Martin et al., 1999, and Vu et al., 1998, in the following) or making explicit the meaning of the target word (cf. Rayner et al., 1994) should contribute to the predictability of that meaning. In addition, a given discourse segment may be predictive of the topic of the next discourse segment (cf. Ariel, 1990, and Experiment 2), although not necessarily of the word selected to represent to it.

Because a predictive context does not necessarily anticipate a certain word, but rather a certain meaning or concept, cloze probability tests (see, e.g., Schwanenflugel, 1991; Vu et al., 1998) are not an adequate measure of concept predictiveness. They focus on the predictability of a specific word from its previous context rather than on the predictability of a concept. Reading times or response times (RTs), however, are. Consider, for instance, the following joke (taken from Coulson & Kutas, 1998):

1. By the time Mary had had her fourteenth child, she'd finally run out of names to call her
   a. Nonjoke ending: offspring
   b. Joke ending: husband

Although *husband* and *offspring* had similar cloze probability (4% and 2%, respectively), *husband* took longer to read than *offspring*. Indeed, although the word *offspring* is quite infrequent and therefore unpredictable, the concept it represents (child) in the given context is not. In contrast, given the same specific context, the concept of "husband" is unpredictable (requiring a "frame shift," to cite Coulson & Kutas). No wonder it took longer to read than *offspring*. Reading times, then, are better adept than cloze probability tests at indexing concept availability and predictability.

The purpose of our study is to lend support to the claim that comprehension involves two separate mechanisms that do not interact initially. Experiment 1 was, therefore, designed to question the direct access view by showing that the selective access attested to by Vu et al. (1998) may very well be a product of the predictive mechanism described previously. Experiment 2 was designed to further provide evidence in favor of the independence of lexical accessing of contextual processes.

EXPERIMENT 1

Experiment 1 aims to show that Vu et al.'s (1998) findings may have an alternative explanation and need not be attributed to early context effects that interact with lexical processes. Rather, they could be due to a mechanism that does not involve interaction with lexical accessing.

In a series of naming experiments, Vu et al. (1998) examined the effect of a strong prior context on the activation of meanings of ambiguous words. They gave participants strong sentential contexts biasing the last target word toward either the salient (dominant) or the less salient (subordinate) meaning. For example, in 2, the context is strongly suggestive of the salient/dominant ("racquet") sense of the last ambiguous word *bat*; in 3, the context is strongly suggestive of the less salient ("mammal") sense of the word:

2. The slugger splintered the *bat.**
   (Probes displayed at *: salient–wooden; unrelated–safe; less salient–fly; unrelated–station).
3. The biologist wounded the *bat.**
   (Probes displayed at *: salient–wooden; unrelated–safe; less salient–fly; unrelated–station).

When participants had read such sentences, they had to name one of four probes (presented in 2 through 3). Findings show that they always named the contextually compatible probe faster than the unrelated one, regardless of salience. On the face of it, then, such findings demonstrate context effects on lexical access: They show that only the contextually appropriate meaning was tapped initially, irrespective of salience. Indeed, if these findings were a result of lexical access, they would question the salience-based view, because in this view, lexical accessing should be sensitive only to linguistic information and to its degree of salience.

To defend the graded salience hypothesis, we attempted to replicate Vu et al.'s (1998) findings by presenting the probes in the sentence pre-final position—that is, immediately *before* the final (target) word is even encountered. We assumed that toward the end of a strong sentential context, the guessing, predictive mechanism is most powerful so that context may avail the compatible meaning before the lexical stimulus is encountered and accessed. We thus predicted that when targets are placed at the end of a strong sentential context, that context would avail the compatible meaning before the lexical stimulus is encountered. Thus, if, under such a condition, Vu et al.'s findings are replicated, these findings would question the interaction hypothesis: They would show that the priming effects obtained by Kellas, Vu, and their colleagues may have been produced by the context alone rather than by accessing words in context (because no accessing was allowed).

In our study, 60 native speakers of English were presented the sentence contexts used by Vu et al. (1998) and had to make lexical decisions as to whether a probe presented before (rather than after) the last (target) word was a word or a nonword.

Method

Design. A 2 × 3 factorial design was used with context type (salient/less salient bias) and probe type (salient/less salient/unrelated) as within-participant factors.

Participants. Sixty participants (32 women and 28 men), ranging from 21 to 60 years old, served as paid participants. They were all native speakers of English. Most of them were students of the Medical School of Tel Aviv University (a special program for North American students), and the rest were North American teachers of English from the Open University of Tel Aviv.

Stimuli. Materials were those used by Vu et al. (1998), which comprised two sentence contexts for 36 homonyms (e.g., *bat*): A context biased toward the salient (racquet) meaning of the homonym (e.g., *The slugger splintered the bat*) and a context biased toward its less salient (mammal) meaning (e.g., *The biologist wounded the bat*). Before reading the final ambiguous target, a probe was displayed, related to either the salient (e.g., *bat–wooden*) or the less salient meaning (e.g., *bat–fly*) of the target word, or was unrelated (e.g., *bat–station*):

4. The slugger splintered the* *bat.*
   (Probes displayed at *: salient–wooden; less salient–fly; unrelated–station).
5. The biologist wounded the* *bat.*
   (Probes displayed at *: salient–wooden; less salient–fly; unrelated–station; manipulated items taken from Vu et al., 1998).[2]

The combination of two context types and three probe types created six conditions. Each participant saw each homonym in only one condition, selected randomly. Thirty-six additional sentences were used as fillers and were always presented with a nonword probe. The 72 items were arranged randomly and displayed in a different order for each participant. Their presentation and response collection were controlled by a Pentium PC, using a C++ program.

[2]It should be noted that, in most cases, placing the probes before the target word did not result in natural, well-formed continuations, as might be deduced from some of the examples.

Procedure. Participants were tested individually. They were first given instructions and had three training trials to make sure they understood the task.

Stimulus presentation resembled a moving window (cf. Katz & Ferretti, 2001/this issue) in which the sentence contexts were displayed word by word across the computer screen at a pace established previously for each participant (see following Pretest section). The words remained visible until the probe was displayed in screen-center position and reappeared after the participant had made a lexical decision as to whether a letter string (the probe) was a word or a nonword in English. The participants responded by pressing one of two (yes/no) keys. The final word of the sentence was then added. In 25% of the cases, a yes/no comprehension question was also displayed. The latency between the onset of the probe and the pressing of the key was measured by the computer and served as RT.

Pretest. To establish the individual pace of presentation of stimuli, each participant read 10 sentences off the computer screen immediately before the actual experiment. The reading time per word was recorded and averaged by the computer and served as the reading pace of the experimental sentences for that participant.

Results

As predicted, because targets were placed at the end of a strong sentential context, that context availed the compatible meaning even *before* the lexical stimulus was encountered and accessed. This was true of both the participant (Fs) and item (Fi) analyses.

We averaged the RT of all trials in each condition. RT outliers above or below two standard deviations from the mean were excluded from the analyses (about 10%). One participant was replaced because he did not respond to the comprehension questions correctly. Means and standard deviations for all conditions are presented in Table 1. Correct responses to word probes (about 96%) were subjected to two-way analyses of variance (ANOVAs). The ANOVAs included two within-participant/items factors: context type (salient/less salient bias) and probe type (salient/less salient/unrelated). Two significant effects were found: a probe-type effect, $F\text{s}(2, 118) = 5.49$, $p < .01$, $F\text{i}(2, 68) = 1.63$, $p = .20$, and, more important, a Context-Type × Probe-Type interaction, $F\text{s}(2, 118) = 15.92$, $p < .0001$, $F\text{i}(2, 68) = 9.18$, $p < .0005$. Specifically, six planned comparisons between means were performed. Within the salient-biased context, there was a significant difference between the RT to the salient and to the less salient probe in the participant analysis, $F\text{s}(1, 59) = 6.65$, $p < .05$, and (marginally so) in the item analysis, $F\text{i}(1, 34) = 3.28$, $p = .079$. The difference between the salient probe and the unrelated probe was also significant, $F\text{s}(1, 59) = 5.44$, $p < .05$, approaching significance in the item analysis, $F\text{i}(1, 34) = 3.77$, $p = .06$, whereas there was no significant difference between the less salient and the unrelated probes, $F\text{s} < 1$, $F\text{i} < 1$.

TABLE 1
Mean Response Times and Standard Deviations (in Msec) to Probes
by Context Type—Experiment 1

	Salient Probe		*Less-Salient Probe*		*Unrelated Probe*	
Context	*M*	*SD*	*M*	*SD*	*M*	*SD*
Salient	951	252	1,003	243	1,005	255
Less-salient	1,057	275	927	237	994	231

A different pattern of results emerged within the context biased toward the less salient interpretation. Again, there was a significant difference between the RT to the salient and to the less salient probes, Fs(1, 59) = 40.67, $p < .0001$, Fi(1, 34) = 18.04, $p < .0005$, but this time in the opposite direction (i.e., the less salient probe was responded to faster than the salient probe). The difference between the salient and the unrelated probes was also significant, Fs(1, 59) = 7.16, $p < .01$, Fi(1, 34) = 6.53, $p < .05$, and so was the difference between the less salient and the unrelated probes, Fs(1, 59) = 12.74, $p < .001$, approaching significance in the item analysis, Fi(1, 34) = 3.70, $p = .06$.

Discussion

Experiment 1, then, replicates Vu et al.'s (1998) findings. Replication of Vu et al.'s findings under conditions that disallow lexical accessing supports our hypothesis that these findings were not necessarily affected by an interactionist mechanism, as assumed by Vu et al. Given that these results were replicated even *before* the target (ambiguous) word was encountered—that is, before lexical accessing could even take place—they cannot be solely attributed to context effect on lexical access. Rather, they show that contextual information was strong enough to predict the appropriate meanings on its own accord. As we assumed, when targets are placed at the end of a strong (sentential) context, contextual information can be strong and effective enough to avail the appropriate interpretation very early on, even before the relevant stimulus is encountered. Our findings thus allow for an alternative explanation to the interactionist hypothesis, which assumes that strong contextual information may result in selective access. (For an alternative critique of Vu et al.'s findings, suggesting that it is the choice of items that is responsible for their results, see Binder & Rayner, 1999.)

EXPERIMENT 2

Experiment 1 demonstrates that the priming results, reported by Vu et al. (1998), are not necessarily a response to the ambiguous word and need not serve as decisive evidence for contextual constraints on lexical access. Instead, we suggest an alter-

native, not less viable account, on which it is another, central, expectation-driven mechanism operating during language comprehension alongside linguistic processes that is responsible for the results obtained by Vu et al. (1998).

One can still argue, however, that the predictive processes assumed by the graded salience hypothesis do not just run in parallel but also eventually penetrate lexical accessing. Although we have shown that context was strong enough to predict the contextually appropriate meaning even before the target word is processed, we have not yet shown that this strong context has no effect on lexical access.

To show that the lexical mechanism is indeed encapsulated with respect to contextual information, it is essential to specify the conditions under which salient but contextually incompatible meanings would not be outweighed by the contextually compatible interpretation. It is our assumption that this expectation-driven mechanism operates most efficiently toward the end rather than at the beginning of sentences. (No wonder we quite often attempt to finish rather than start a sentence for a slow speaker.) We therefore predict that even a strong context would not inhibit salient meanings at the beginning of sentences. This prediction is inconsistent with those of the interactive models, which assume that, in a rich and supportive context, the appropriate meaning is tapped initially, directly, and exclusively, without involving contextually incompatible meanings at all, because the latter are inhibited by contextual information (cf. Martin et al., 1999; Vu et al., 1998; see also Glucksberg, Newsome, & Goldvarg, 2001/this issue).[3]

To tease apart the two mechanisms, we used novel, one-word metaphors. Such metaphors have a coded, contextually incompatible meaning (the salient, literal meaning) and an uncoded contextually appropriate interpretation (the nonsalient referential interpretation). Priming effects related to the salient, contextually incompatible meaning can thus only be attributed to the lexical mechanism. In contrast, priming effects

[3]Glucksberg, Newsome, & Goldvarg (2001/this issue) suggested that contextual information inhibits literal meanings irrelevant to metaphor interpretation. For instance, following *"My lawyer was a shark,"* a metaphorically irrelevant literal sentence probe *Geese can swim* was read more slowly relative to a metaphorically relevant literal sentence probe *Geese are vicious* (which contains a metaphor relevant property "vicious"). It is possible, however, that the irrelevant shark property "swimming" is also less salient compared to "viciousness," and hence less accessible, as can be deduced from Glucksberg et al.'s own findings. In their study, the probe sentence substantiating "viciousness" was read faster than the probe sentence substantiating "swimming," regardless of prime. Thus, *Geese are vicious* took almost equally long to read (1,578 msec vs. 1,568 msec) whether the prime was metaphoric (*"My lawyer is a shark"*) or literal ("This hammerhead is a shark") but faster than *Geese can swim*, whether preceded by a literal (1,701 msec) or a metaphoric (1,926 msec) prime. Such findings suggest that there was no inhibition of irrelevant properties: After all, *Geese can swim* is relevant in the context of the literal prime ("This hammerhead is a shark") in which "swimming" is relevant. Still, it was read more slowly than the relevant *Geese are vicious*. Relevance, then, cannot account for all the findings. Instead, the salience-based explanation seems to account for all the findings: It is low salience rather than low relevance probes that took longer to read than high salience rather than high relevance probes. (It should be noted that Glucksberg et al. did not find priming for any property in a word–word priming test, which indeed often fails to show differences.)

related to the nonsalient, contextually compatible interpretation can only be attributed to the expectation-driven mechanism (see also Gerrig, 1989).[4]

To test our hypothesis, we conducted Experiment 2, this time in Hebrew, employing novel metaphors whose order of presentation was manipulated. We predicted that, at the beginning of sentences, context effects would not outweigh salience effects. Moreover, even a strongly biased context would not inhibit salient meanings: Salient "false positives" would be activated, regardless of contextual information to the contrary. Experiment 2, then, compares context effects in sentence initial and final position vis à vis salience effects. The contexts were contrived in such a way as to be predictive of the next topic, which was also the target word.

To establish the salience of the meanings out of a biasing context, we first conducted a pretest that measured RTs to the experimental probes when the targets were embedded in a neutral context, thus annulling context effects and tapping degree of salience exclusively.

Pretest

Method

Design. A simple design was used, with only one within-participant factor—probe type (salient/nonsalient/unrelated).

Participants. Twenty-four undergraduate and graduate students (16 women and 8 men) of Tel Aviv University, ranging from 20 to 40 years old, served as paid participants. They were all native speakers of Hebrew.

Stimuli. Twenty-four sentences were used containing the 24 targets (e.g., *delinquents*) to be used metaphorically (referring to "kids") in the experimental sentences. The words preceding the targets made up unbiased, neutral contexts:

6. Neutral context: This place is full of delinquents.*
   (Probes displayed at *: salient–criminals; nonsalient–kids; unrelated–painters).

Immediately (0 msec) after offset of the final target word (*delinquents*), a probe was displayed (at *, see Example 6) related to either the salient meaning of the target (corresponding to its literal meaning, e.g., "criminals") or the nonsalient meaning (corresponding to its metaphorical interpretation, e.g., "kids"), or was unrelated to any of the senses of the target (e.g., "painters"). Each participant saw each sentence

[4]Several researchers proposed the parallel process to account for comprehension of nonliteral language; see, for instance, Dews and Winner (1999); Keysar (1989); and Ortony, Schallert, Reynolds, and Antos (1978).

only once, followed by one of the three probes. Twenty-four additional sentences were created and served as fillers. They were always presented with a nonword probe. The 48 items were arranged randomly and presented in a different order for each participant. All stimuli were presented in Hebrew. Their presentation and response collection were controlled by a Pentium PC, using a C++ program.

Procedure. The procedure was the same as used in Experiment 1.

Results and Discussion

We used the same averaging and exclusion of outliers procedure as in Experiment 1. RT outliers were excluded from the analyses (about 8%). Means and standard deviations of RTs for the correct responses (about 96%) for the three conditions are presented in Table 2. To see whether the salient probe is accessed more rapidly than the nonsalient and unrelated probes as assumed, we used a planned comparison of the means, with the contrast (–1, .5, .5) for the salient, nonsalient, and unrelated probes, respectively. The contrast was significant for both, the participant—*F*s(1, 23) = 6.13, $p < .05$—and item—*F*i(1, 23) = 14.57, $p < .001$—analyses, confirming the difference of availability predicted for salient and nonsalient meanings.

Experiment

Given the salience difference between the various interpretations of the target words, Experiment 2 was designed to show that effects of a strong prior context would not override salience effects in sentence initial position. Notwithstanding these effects, salient meanings would not be blocked, not even where context is expected to be superior (i.e., at the end of sentences).

Specifically, based on our assumption that context cannot interact with lexical accessing, we predict that probes related to the salient, contextually incompatible (literal) meaning of the metaphors will be primed, regardless of prior context and sentential position. In addition, probes related to the nonsalient, contextually com-

TABLE 2
Mean Response Times and Standard Deviations (in Msec) to Probes in the Pretest of Experiment 2

	Salient Probe	*Nonsalient Probe*	*Unrelated Probe*
M	710	747	752
SD	172	175	159

patible (metaphoric) interpretation will be primed, too, due to the effects of the expectation driven mechanism. These effects, however, will not override salience effects in sentence initial position, but may be faster than salience effects in noninitial position (cf. Experiment 1).

Method

Design. A 2 × 3 factorial design was used with probe position (initial/final) and probe type (salient/nonsalient/unrelated) as within-participant/item factors.

Participants. Sixty undergraduate and graduate students (43 women and 17 men) of Tel Aviv University, ranging from 19 to 32 years old, served as paid participants. They were all native speakers of Hebrew.

Stimuli. The same target words and probes used in the pretest were used here as well. For each target word (*delinquents*), a sentence was comprised so that the target appeared in either sentence initial (7) or sentence final (8) position. (See Appendix for sample examples.) A short passage was constructed, strongly biasing the target sentence toward the nonsalient (metaphorical) meaning, which in all cases was the topic of the previous context sentence as well as the topic of target sentence, creating strong expectations (cf. Ariel, 1990; Gernsbacher, 1990; Giora, 1985a, 1985b; Reinhart, 1980). Consequently, the targets were ambiguous between the contextually compatible but nonsalient (metaphoric) meaning and the salient but contextually incompatible (literal) meaning:

7. Ambiguous–initial context: Sarit's sons and mine went on fighting continuously. Sarit said to me: These delinquents* won't let us have a moment of peace. (Probes displayed at *: salient–criminals; contextually compatible–kids; unrelated–painters)
8. Ambiguous–final context: Sarit's sons and mine went on fighting continuously. Sarit said to me: A moment of peace won't let us have these delinquents*.[5] (Probes displayed at *: salient–criminals; contextually compatible–kids; unrelated–painters)

The combination of two probe positions and three probe types created six conditions. Each participant saw each passage in only one condition selected ran-

[5]The word ordering in Hebrew is such that the target NP occupies initial position, preceding the demonstrative: The *delinquents* these won't let us have a moment of peace. A moment of peace won't let us have the *delinquents* these.

domly. As earlier, 24 additional sentences were created as fillers and presented with a nonword probe. The 48 items were arranged randomly and presented in a different order for each participant.

Procedure. As in Experiment 1, participants read the passages at their own natural reading pace established in a pretest (cf. Experiment 1). They had to make a lexical decision as to whether the probe presented immediately (0 msec) after offset of the target word was a word or a nonword.

Results and Discussion

We used the same averaging and exclusion of outliers procedure (excluding about 6%) as in Experiment 1. One participant was replaced because he did not respond to the comprehension questions correctly. Means and standard deviations of RTs for the correct responses (about 98%) for the six conditions are presented in Table 3 and illustrated by Figure 1. The analyses of variance included two within-participant/item factors: target position (initial/final) and probe type (salient/nonsalient/unrelated). We conducted four planned comparisons between means. In the initial position, there was a significant difference between the RT to the salient and to the unrelated probe, Fs(1, 59) = 22.92, $p < .0001$, approaching significance in the item analysis, Fi(1, 23) = 4.04, $p = .056$. The difference between the nonsalient probe and the unrelated probe was also significant, Fs(1, 59) = 14.50, $p < .0005$, Fi(1, 23) = 6.39, $p < .05$. In the final position, the difference between the salient and the unrelated probes approached significance (in the participant analysis) Fs(1, 59) = 3.94, $p = .052$, Fi(1, 23) = 1.20, $p = .29$. However, the difference between the nonsalient and the unrelated probe was significant (in the participant analysis), Fs(1, 59) = 9.56, $p < .005$, Fi(1, 23) < 1.[6]

Results are indeed consistent with our predictions, disconfirming the inhibitory assumption of the direct access view, while supporting the graded salience, reordered access, and modular views. They show that although contextual information availed the appropriate meaning in sentence initial as well as in sentence final position, it did not mute salient but contextually incompatible meanings. They further show that, as anticipated, in sentence final position contextual effects were somewhat faster than salience effects, emerging probably before the target word was encountered and processed (cf. Experiment 1). Of importance, however, these effects did not inhibit salient though inappropriate meanings. False positives, then,

[6]The RTs at sentence final position are slightly inflated. This, however, might be due to sentence wrap-up effects, effects occurring at the end of the sentence.

TABLE 3
Mean Response Times and Standard Deviations (in Msec) to Probes by Target Position in Experiment 2

	Salient (Contextually Incompatible) Probe		*Nonsalient (Contextually Compatible) Probe*		*Unrelated Probe*	
Target Position	*M*	*SD*	*M*	*SD*	*M*	*SD*
Initial position	1,021	217	1,019	241	1,097	247
Final position	1,088	282	1,046	240	1,123	273

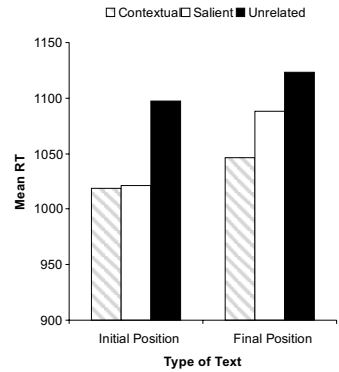

FIGURE 1 Mean RTs (in msec) to probes related to the salient (contextually incompatible) and nonsalient (contextually compatible) meanings of the target words and unrelated probes in Experiment 2.

were activated on account of their saliency, regardless of contextual compatibility (see Figure 1).

Such findings support our view that language comprehension involves two distinct mechanisms that run in parallel, one sensitive to contextual information and one sensitive to coded, salient information. Thus, although contextual information may have fast effects, they do not filter out salience effects. Salient meanings are activated on encounter of the verbal stimulus, irrespective of context predictiveness.

GENERAL DISCUSSION

This study tested the assumption that early comprehension processes involve two separate mechanisms that operate simultaneously without interacting. Lexical accessing is modular and stimulus driven; it is a bottom-up perceptual mechanism induced by a lexical stimulus to search the mental lexicon for its match. This mechanism is encapsulated with respect to nonlinguistic information and thus operates locally (i.e., on the word level). Lexical access is exhaustive and ordered: Salient meanings are activated first.

Contextual effects, on the other hand, are the outcome of a central, expectation-driven mechanism that operates globally during language comprehension at the point where prior linguistic information has already been processed and interfaced with other cognitive processes (e.g., inferencing). It does not affect lexical accessing and hence does not prevent contextually incompatible meanings.

Experiment 1 indicates that contextual facilitation of the compatible meaning of a target word can occur even before that target is encountered—that is, before lexical accessing takes place, fostering an impression of a selective process. Experiment 2 shows that although context effects may be fast, they do not inhibit salient but contextually incompatible meanings, not even where context may be most effective (i.e., in sentence final position). Of importance, sentence ordering (initial vs. noninitial) may be crucial for the operation of the predictive mechanism. Thus, we would not expect contextual effects to temporally outweigh salient meanings in the beginning of sentences. In addition, context effects may enable uncoded, novel interpretations to become immediately available, even before lexical accessing is afforded, both in sentence final as well as in sentence initial position (Experiment 2).

Assuming two different mechanisms as opposed to a single, interactionist mechanism may not just account for our findings, but may better account for conflicting findings prevalent in the literature. Thus, findings demonstrating that a strong context can avail the appropriate meaning immediately, regardless of salience (e.g., Vu et al., 1998), can also be viewed as induced by a contextual mechanism per se, one that does not interact with lexical processing. Indeed, in Vu et al., in which targets were placed in sentence final position, context effects temporally preceded lexical accessing, suggesting that findings compatible with an interactionist account may very well be the product of contextual processes that do not interact with lexical accessing.

In contrast, findings showing that contextually incompatible meanings slow down processes may be due to the lexical mechanism, particularly if targets are placed in sentence initial position. In spite of a strong prior context, context effects in such a position are expected to neither inhibit nor supercede salient though contextually incompatible meanings. In some cases, they may also lag behind lexical accessing. Hence, the predicted interference of salient but incompatible meanings in the interpretation process of targets placed in initial position. Indeed, in Gibbs (1990), metaphors such as 9 took longer to read than literal equivalents such as 10, irrespective of the prior story contexts, which were rich and supportive. Placed in sentence initial position, targets (*creampuff/fighter*) were probably accessed initially literally via their salient meaning. In 9, such processing resulted in contextual misfit, inducing longer reading times necessary for the resolution of the conflict. In 10, however, context and salience coincided. Hence, no extra processing was required.

9. The *creampuff* didn't even show up. (taken from Gibbs, 1990)

10. The *fighter* didn't even show up. (taken from Gibbs, 1990)

However, when critical words (*creampuff/loser*) were placed in clause final position, as in 11 and 12, the differences noted disappeared (cf. Onishi & Murphy, 1993), as predicted:

11. He's such a *creampuff* that he didn't even show up, said Tracey. (taken from Onishi & Murphy, 1993)
12. He's such a *loser* that he didn't even show up, said Tracey. (taken from Onishi & Murphy, 1993)

In sum, our studies show that initial lexical processes are independent of contextual processes. Although context may have early effects occurring even before lexical accessing takes place (see also Rayner, Binder, & Duffy, 1999), they do not affect lexical accessing and therefore do not block salient meanings. False positives are accessed on account of their salience, regardless of contextual information to the contrary.

Our studies further show that, along the lines suggested by Fodor (1983), comprehension involves an additional, expectation driven mechanism, which allows for the processing of novel, nonsalient interpretations. In our studies (cf. Experiment 2), uncoded, nonsalient, but contextually compatible interpretations were available immediately even in sentence initial position (following a strong prior context). Such findings can only be due to a contextual mechanism that operates globally without interacting with lexical processes (cf. Experiment 1). Indeed, these findings cannot be explained by most of the existing models, which are adept at dealing with coded (salient and less salient) meanings only, while assuming an interactionist mechanism (e.g., the direct access view, the reordered access view). The graded salience hypothesis, then, seems better than the existing models at explaining the diversity of findings abounding in the literature (see also Giora, in press-a, in press-b, for a salience-based account of literal and figurative language comprehension).

ACKNOWLEDGMENTS

Support for this research was provided by The Graduate Program for Cognitive Studies of Language and Its Uses, The Porter School of Cultural Studies, Tel Aviv University to Orna Peleg, and by Grant 891/96–7 from The Israel Science Foundation to the Rachel Giora.

We also thank Olga Zeigelman for preparing the software and providing statistical analyses. We are particularly indebted to Mira Ariel, Steven Frisson, Vu Hoang, George Kellas, Albert Katz, and Keith Rayner for their highly valuable comments.

REFERENCES

Ariel, M. (1990). *Accessing NP antecedents.* London: Routledge.

Binder, K. S., & Rayner, K. (1998). Contextual strength does not modulate the subordinate effect: Evidence from eye fixations and self-paced reading. *Psychonomic Bulleting & Review, 5,* 217–276.

Binder, K. S., & Rayner, K. (1999). Does contextual strength modulate the subordinate effect? A reply to Kellas and Vu. *Psychonomic Bulleting & Review, 6,* 518–522.

Carston, R. (1999, June). *Linguistic meaning and literal meaning.* Paper presented at Pragma99: International pragmatics conference on pragmatics and negotiations, Tel Aviv University, Tel Aviv, Israel.

Connine, C. M., Blasko, D., & Wang, J. (1994). Vertical similarity in spoken word recognition: Multiple lexical activation, individual differences, and the role of sentence context. *Perception & Psychophysics, 56,* 624–636.

Coulson, S., & Kutas, M. (1998). Frame-shifting and sentential integration. Manuscript submitted for publication.

Dews, S., & Winner, E. (1999). Obligatory processing of the literal and nonliteral meanings of ironic utterances. *Journal of Pragmatics, 31,* 1579–1599.

Duffy, S. A., Morris, R. K., & Rayner, K. (1988). Lexical ambiguity and fixation times in reading. *Journal of Memory and Language, 27,* 429–446.

Fodor, J. (1983). *The modularity of mind.* Cambridge, MA: MIT Press.

Frisson, S., & Pickering, M. J. (2001/this issue). Obtaining a figurative interpretation of a word: Support for underspecification. *Metaphor and Symbol, 16,* 149–171.

Gernsbacher, M. A. (1990). *Language comprehension as structure building.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Gerrig, R. J. (1989). The time course of sense creation. *Memory & Cognition, 17,* 194–207.

Gibbs, R. W., Jr. (1980). Spilling the beans on understanding and memory for idioms in conversation. *Memory & Cognition, 8,* 449–456.

Gibbs, R. W., Jr. (1990). Comprehending figurative referential descriptions. *Journal of Experimental Psychology: Learning, Memory and Cognition, 16,* 56–66.

Gibbs, R. W., Jr. (1994). *The poetics of mind.* Cambridge, England: Cambridge University Press.

Giora, R. (1985a). A text-based analysis of nonnarrative texts. *Theoretical Linguistics, 12,* 115–135.

Giora, R. (1985b). Towards a theory of coherence. *Poetics Today, 6,* 699–716.

Giora, R. (1997). Understanding figurative and literal language: The graded salience hypothesis. *Cognitive Linguistics, 7,* 183–206.

Giora, R. (1999). On the priority of salient meanings: Studies of literal and figurative language. *Journal of Pragmatics, 31,* 919–929.

Giora, R. (in press-a). Literal vs. figurative language: Different or equal? *Journal of Pragmatics.*

Giora, R. (in press-b). *On our mind: Salience, context, and figurative language.* New York: Oxford University Press.

Giora, R., & Fein, O. (1999). Irony: Context and salience. *Metaphor and Symbol, 14,* 241–257.

Giora, R., Fein, O., & Schwartz, T. (1998). Irony: Graded salience and indirect negation. *Metaphor and Symbol, 13,* 83–101.

Glucksberg, S., Newsome, M. R., & Goldvarg, Y. (2001/this issue). Inhibition of the literal: Filtering metaphor-irrelevant information during metaphor comprehension. *Metaphor and Symbol, 16,* 277–293.

Katz, A. N., & Ferretti, T. R. (2001/this issue). Moment-by-moment reading of proverbs in literal and nonliteral contexts. *Metaphor and Symbol, 16,* 193–221.

Kawamoto, A. H. (1993). Non-linear dynamics in the resolution of lexical ambiguity: A parallel distributed processing account. *Journal of Memory and Language, 32,* 474–516.

Keenan, J. (1978). Inferring causal connections in prose comprehension. Paper presented at the American Psychological Association Convention.

Keysar, B. (1989). On the functional equivalence of literal and metaphorical interpretations in discourse. *Journal of Memory and Language, 28,* 375–385.

Klin, C. M., Murray, J. D., Levy, W. H., & Guzmán, A. E. (1999). Forward inferences: From activation to long term memory. *Discourse Processes, 27,* 241–260.

Martin, C., Vu, H., Kellas, G., & Metcalf, K. (1999). Strength of discourse context as a determinant of the subordinate bias effect. *The Quarterly Journal of Experimental Psychology, 52A,* 813–839.

Neely, J. (1977). Semantic priming and retrieval from lexical memory. *Journal of Experimental Psychology: General, 196,* 226–254.

Onifer, W., & Swinney, D. A. (1981). Accessing lexical ambiguities during sentence comprehension: Effects of frequency of meaning and contextual bias. *Memory & Cognition, 9,* 225–236.

Onishi, K. H., & Murphy, G. L. (1993). Metaphoric reference: When metaphors are not understood as easily as literal expressions. *Memory & Cognition, 21,* 763–772.

Ortony, A., Schallert, D. L., Reynolds, R. E., & Antos, S. J. (1978). Interpreting metaphors and idioms: Some effects of context on comprehension. *Journal of Verbal Learning and Verbal Behavior, 17,* 465–477.

Peleg, O. (2000). *Linguistic and nonlinguistic mechanisms in language comprehension.* Unpublished doctoral dissertation, Tel Aviv University, Tel Aviv, Israel.

Rayner, K., Binder, K. S., & Duffy, S. (1999). Contextual strength and subordinate bias effect. *Quarterly Journal of Experimental Psychology, 52A,* 841–852.

Rayner, K., & Frazier, L. (1989). Selection mechanisms in reading lexically ambiguous words. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 779–790.

Rayner, K., & Morris, R. K. (1991). Comprehension processes in reading ambiguous sentences: Reflections from eye movements. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 175–198). Amsterdam: North Holland.

Rayner, K., Pacht, J. M., & Duffy, S. A. (1994). Effects of prior encounter and global discourse bias on the processing of lexically ambiguous words: Evidence from eye fixations. *Journal of Memory and Language, 33,* 527–544.

Reinhart, T. (1980). Conditions for text coherence. *Poetics Today, 1,* 161–180.

Schwanenflugel, P. J. (1991). Contextual constraint and lexical processing. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 23–45). Amsterdam: North Holland.

Sereno, C. S., Pacht, J. M., & Rayner, K. (1992). The effect of meaning frequency on processing lexically ambiguous words: Evidence from eye fixations. *Psychological Science, 3,* 269–300.

Sperber, D., & Wilson, D. (1986/1995). *Relevance: Communication and cognition.* Oxford, England: Blackwell.

Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior, 18,* 645–659.

Till, R. E., Mross, E. F., & Kintsch, W. (1988). Time course of priming for associate and inference words in a discourse context. *Journal of Verbal Learning and Verbal Behavior, 16,* 283–298.

Vu, H., Kellas, G., Metcalf, K., & Herman, R. (2000). The influence of global discourse on lexical ambiguity resolution. *Memory & Cognition, 28,* 236–252.

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

APPENDIX
Examples of Test Stimuli

Translated sample items: (a) ambiguous–initial context; (b) ambiguous–final context:

1. Mira and I ate in a restaurant. We ordered steaks, which were terrible. Afterwards, I asked Mira if she wanted us to order a dessert. Mira answered:
   (a) "This rubber* has done away with my appetite."
   (b) "Has done away with my appetite this rubber."*
   (Probes displayed at *: salient–flexible; contextually compatible–meat; unrelated–breakable)

2. Lea, my good friend, works all day without taking a break. Yesterday I spoke with my husband about her and I said:
   (a) "This ant* never takes a day off."
   (b) "Never takes a day off this ant."*
   (Probes displayed at *: salient–fly; contextually compatible–girl; unrelated–elevator)

3. Amir and I were watching a soccer game between Israel and Cyprus when Mizrachi (a player) missed a perfect opportunity to score. Amir told me:
   (a) "This zero* (i.e., this good-for-nothing) disappoints me every time I see him."
   (b) "Every time I see him he disappoints me this zero."*
   (Probes displayed at *: salient–nothing; contextually compatible–player; unrelated–tallith [praying shawl])

4. Dana, Ruth's daughter, is only 14 years old and already 1.82 meters [tall, equivalent to almost six feet tall]. Yesterday while I was visiting them Dana was talking endlessly on the phone. Ruth said:
   (a) "This giraffe* doesn't leave the receiver for a second."
   (b) "Doesn't leave the receiver for a second this giraffe."*
   (Probes displayed at *: salient–animal; contextually compatible–girl; unrelated–sign)

5. Na'ama loves animals and treats her own with as much care as possible. Yesterday she was feeding one of her kittens when I came by. Na'ama said:
   (a) "This princess* eats only imported cheese."
   (b) "Only imported cheese eats this princess."*
   (Probes displayed at *: salient–queen; contextually compatible–cat; unrelated–delegate)

6. Tamara and I met Yossi at the concert and he hardly paid any attention to us. Tamara said:
   (a) "To this iceberg* I almost got married."
   (b) "I almost got married to this iceberg."*
   (Probes displayed at *: salient–freezing; contextually compatible–guy; unrelated–sequence)