

Memory Operations That Support Language Comprehension: Evidence From Verb-Phrase Ellipsis

Andrea E. Martin and Brian McElree
New York University

Comprehension of verb-phrase ellipsis (VPE) requires reevaluation of recently processed constituents, which often necessitates retrieval of information about the elided constituent from memory. A. E. Martin and B. McElree (2008) argued that representations formed during comprehension are content addressable and that VPE antecedents are retrieved from memory via a cue-dependent direct-access pointer rather than via a search process. This hypothesis was further tested by manipulating the location of interfering material—either before the onset of the antecedent (proactive interference; PI) or intervening between antecedent and ellipsis site (retroactive interference; RI). The speed–accuracy tradeoff procedure was used to measure the time course of VPE processing. The location of the interfering material affected VPE comprehension accuracy: RI conditions engendered lower accuracy than PI conditions. Crucially, location did not affect the speed of processing VPE, which is inconsistent with both forward and backward search mechanisms. The observed time-course profiles are consistent with the hypothesis that VPE antecedents are retrieved via a cue-dependent direct-access operation.

Keywords: sentence processing, verb-phrase ellipsis, retrieval interference, speed–accuracy tradeoff, proactive interference

Successful language comprehension routinely requires establishing dependencies spanning several words, phrases, or even sentences. As such, comprehenders must access products of past language processing in order to incorporate the contents of current processing into an evolving discourse. A central issue for models of comprehension thus concerns the nature of memory representations formed during comprehension and how those representations are accessed in real time.

Basic memory research has identified two distinct ways that memories can be retrieved. One way involves a search operation where memory representations are sequentially sampled until the required information is recovered. This type of serial operation appears to be the primary means by which relational (temporal and spatial order) information is recovered (Gronlund, Edwards, & Ohrt, 1997; McElree, 2001, 2006; McElree & Doshier, 1993). The second way presupposes content-addressable memory representations, so that information in the retrieval context (cues) provides direct access to relevant memory representations, without searching through extraneous representations (McElree & Doshier, 1989). Cue-dependent direct-access operations can be implemented in memory models with diverse storage architectures (Clark & Gron-

lund, 1996) and have been found to be the primary means through which item representations are accessed (e.g., McElree, 2000, 2006; McElree & Doshier, 1989, 1993).

Which type of operation subserves language comprehension? Much of the information required for comprehension is relational, which, based on memory research with verbal lists (Gronlund et al., 1997; McElree, 2001, 2006; McElree & Doshier, 1993), seems to require a search process. For example, hierarchical sentence structure is often encoded by constituent order in languages like English (McElree, Foraker, & Dyer, 2003). If a dependency necessitates a constituent in memory with particular morphology or requires a constituent in a specific syntactic/semantic role associated with a sentence position, then comprehenders might need to serially search through the input for the required constituent (McElree et al., 2003). At one extreme, the input representation might be scanned either forward or backward, with each constituent evaluated sequentially for its degree of match to the search criteria. Alternatively, the search might operate over more structured representations, with the candidate set restricted to those that match content-based criteria.

The key prediction of a search operation is that retrieval speed will decrease as more elements are added to the candidate set. Serial models predict that retrieval speed decreases linearly with the number of elements searched, whereas parallel models often predict nonlinear decreases (McElree & Doshier, 1989; Townsend & Ashby, 1983). In contrast, additional elements in memory need not affect retrieval speed in a content-addressable system with a direct-access operation. Additional material in memory might reduce the likelihood that the target item is successfully retrieved from memory. If the additional material overlaps in content with the target constituent, it may engender interference at encoding or during storage, resulting in a reduction of distinctiveness of the stored representation of the target item that,

Andrea E. Martin and Brian McElree, Department of Psychology, New York University.

This research was supported by a National Institutes of Health grant (R01-HD056200) awarded to Brian McElree and a National Science Foundation Graduate Research Fellowship (2006025605) awarded to Andrea E. Martin.

We thank TS for helpful commentary on earlier versions of this article.

Correspondence concerning this article should be addressed to Andrea E. Martin, Department of Psychology, New York University, 6 Washington Place, 8th Floor, New York, NY 10003. E-mail: andrea.martin@nyu.edu

in turn, reduces the accuracy of retrieval (McElree et al., 2003; Van Dyke & McElree, 2006). Additional material could also engender interference at retrieval if the contents of that material partially match the cues that are used at retrieval, producing what has been referred to as *cue-overload*, where retrieval cues are no longer sufficient to guarantee successful recovery of the target and to exclude the recovery of inappropriate items (Nairne, 2002; Öztekin & McElree, 2007; Watkins & Watkins, 1975). However, a direct-access operation enables representations of differing quality and/or distinctiveness to be retrieved in equal time (e.g., McElree & Doshier, 1989; Ratcliff, 1978). As such, measures of the speed with which a constituent is retrieved during comprehension as a function of the number of elements in memory provide one means of investigating which type of retrieval operation is used in comprehension.

Retrieval of Predicted and Unpredicted Dependencies

McElree et al. (2003; see also McElree, 2000) used this logic to investigate the retrieval operations used to resolve subject–verb dependencies and filler–gap dependencies (e.g., in clefts such as *This was the book that the editor admired*, where *the book* must be associated as the direct object of the final verb, *admired*). Different amounts of material were interpolated between the subject and verb in the first case or between the filler and gap in the second. Increasing the amount of interpolated material reduced the probability of computing an acceptable interpretation, reflected in lower accuracy for acceptability judgments (see below) but not comprehension speed. This suggests that linguistic representations are content addressable, directly accessible via the cues provided by the matrix verb.

One potential concern about these studies is that they used dependencies with constituents the grammar marks as having a future role in the sentence—subjects must be unified with verbs, and fillers must receive a grammatical and semantic role from the gap position. It is possible, therefore, that these constituents were assigned some special status in memory. In fact, models of parsing often assume that these types of constituents are held in specialized stacks or buffers, which can mimic direct access (McElree et al., 2003).

However, Martin and McElree (2008) used similar logic to investigate verb–phrase ellipsis (VPE) processing. Crucially, antecedents of VPE are fully integrated in local context; there is no grammatical marking signaling a further role downstream. Without special status, VPE engenders retrieval in situ, triggered only by cues at the ellipsis site.

Martin and McElree (2008) varied distance between antecedent and VPE, and antecedent complexity. Neither affected retrieval speed, but as in subject–verb and filler–gap dependencies, interpreting more material within the dependency decreased accuracy. Collectively, these studies support the hypothesis that the antecedent representation is content addressable and retrieved with a cue-dependent direct-access operation. To our knowledge, no other extant theory explicitly states how the antecedent of VPE is recovered from memory, although some have assumed a search process (e.g., Murphy, 1985).

Evaluating Alternative Search Mechanisms

Extant time-course studies of nonadjacent dependencies are consistent with direct-access retrieval and inconsistent with a large

class of search mechanisms. They are incompatible with serial and parallel exhaustive searches, as the number of elements in memory determines retrieval speed under these mechanisms (see McElree & Doshier, 1989; McElree et al., 2003). Likewise, they are inconsistent with the backward (self-terminating) search found to underlie the recovery of relational information, as material interpolated within the dependency is predicted to slow retrieval. However, the results are not necessarily inconsistent with a forward serial search, where the comprehender starts at the beginning of a sentence or discourse and searches forward for a constituent to resolve the dependency. Forward search could be viewed as more plausible than backward search given findings that forward recall is less taxing than backward recall (e.g., Thomas, Milner, & Haberlandt, 2003) and comprehension advantages found for first-mentioned material (e.g., Gernsbacher & Hargreaves, 1988). Crucially, no effect of interpolated material is predicted by forward search if the target item is sentence initial, as it was in McElree (2000), McElree et al. (2003), and Martin and McElree (2008).

To investigate whether forward search underlies the retrieval of VPE antecedents, we contrasted conditions in which additional material occurred either before the antecedent—which creates proactive interference (PI)—or between the antecedent and VPE—which creates retroactive interference (RI). If forward search mediates access to the antecedent, then processing time at the VPE should be slower under PI than RI. Backward search predicts the opposite pattern, with processing being slower under RI than PI. In contrast, no difference in processing time is predicted for a direct-access operation, as cues at retrieval make direct contact with relevant representations by virtue of their content.

Measuring Processing Speed

Differences in processing speed provide the crucial data for discriminating between retrieval operations in VPE. Nonetheless, as discussed earlier, the two interference conditions may also differentially affect processing accuracy. Differences in accuracy are less diagnostic as to the type of retrieval operation, because they can arise from factors that are orthogonal to the nature of the specific retrieval mechanism and that could affect either search or direct access. For example, the additional material is more recent in RI compared with PI conditions, and so it may provide a stronger source of retrieval interference, due either to differences in representational strength of the antecedent or in cue diagnosticity, both of which, in turn, would negatively affect the probability of successful retrieval. Thus, although accuracy, under either forward or backward search, should mirror processing speed (viz., accuracy should be lower under RI than PI under backward search, whereas the reverse pattern is predicted under forward search; see McElree & Doshier, 1993), potential accuracy differences could simply reflect interference-contingent differences in these other factors.

Identifying the types of material in memory that engender interference at retrieval provides insights into the nature of the retrieval operation, irrespective of whether it is a search or direct-access mechanism. Under a search mechanism, identifying properties of the material that slow or disrupt the search process provides insights into the criteria and decision processes that are used to select items recovered in the search. In a cue-driven direct-access mechanism, properties of the interfering material

provide insight into what particular cues drive retrieval. To explore this issue, we used interpolated material in both the RI and PI conditions that varied in the degree to which it overlapped with the target in morphosyntactic, semantic, and/or pragmatic properties. Specifically, we varied whether the interpolated material contained temporal adjuncts or additional verb phrases, as well as whether the interpolated verb phrase was headed by a verb that matched the cue for the elided target. We did this by contrasting (a) ellipsis cued by “did” and an interpolated verb phrase including “was” and vice versa (a partial match between cue and interpolated material) with (b) ellipsis cued by “did” or “was” and an interpolated verb phrase with a matching “did” or “was,” respectively (a full match between cue and interpolated material).

Because interfering information can potentially affect speed and accuracy, a behavioral method that conjointly measures speed and accuracy is required to discriminate between retrieval mechanisms. The response-signal speed-accuracy tradeoff (SAT) procedure is such a method (e.g., Doshier, 1979; Reed, 1973; Wickelgren, 1977). We used the multiple response variant of the procedure (e.g., Martin & McElree, 2008; McElree, 1993), where participants read sentences phrase by phrase and then decided, at the VPE,

whether it was acceptable. Participants were trained to respond to a signal (tone) presented 14 times at 350-ms intervals following the onset of the VPE. The sampled times enabled us to fully measure how the interpretations of VPE unfolded over time.

Figure 1 illustrates SAT functions— d' accuracy versus processing time—for two hypothetical conditions. Characteristic functions show a period of chance performance ($d' = 0$), a period of increasing accuracy, and an asymptotic period where further processing does not improve performance. The form of the functions can be summarized with an exponential approach to a limit (Equation 1), which quantifies how the interpretation of VPE unfolded over time:

$$d' = \lambda[1 - e^{-\beta(t-\delta)}] \text{ for } t > \delta, \text{ otherwise } t = 0. \quad (1)$$

The asymptotic parameter (λ) reflects the accuracy reached with maximum processing time. Figure 1A illustrates a pure asymptotic difference. Such differences may arise if additional material affects the probability of retrieving the correct antecedent. The crucial predictions concern the β and δ parameters, which estimate the function's dynamics (viz., how quickly accuracy accrues to its asymptote). Figure 1B illustrates a difference in rate and intercept.

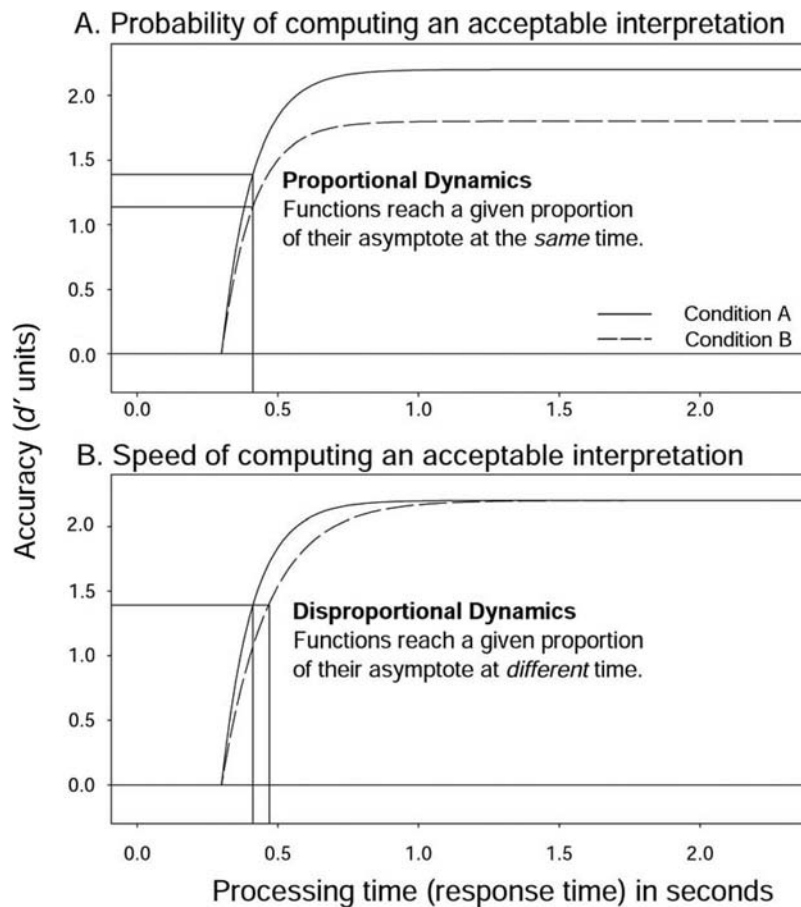


Figure 1. Hypothetical speed-accuracy tradeoff functions illustrating two conditions that differ by asymptote only (A) or rate (B). The intersection of the horizontal and vertical lines shows the point in time (abscissa) when the functions reach two thirds of their respective asymptote (ordinate). When dynamics are proportional (A), the functions reach the two-thirds point at the same time.

The parameter δ estimates the function's intercept, when participants are first able to discriminate acceptable from unacceptable. The parameter β estimates the rate at which accuracy grows from chance to asymptote. Differences in some combination of these two parameters are expected if accessing the antecedent requires a search. For example, under forward search, processing material prior to the antecedent should slow interpretation of VPE relative to processing material after the antecedent, delaying the intercept (δ) or reducing the rate (β) of the SAT function.

Method

Participants

Twelve native speakers of American English (aged 18–26 years) from the New York University community were paid \$10/hour to participate. They completed nine 1-hr sessions and a 45-min practice session.

Materials

Thirty-six sets, illustrated in Tables 1 and 2, were used. The main contrast concerned VPE with material before the antecedent and the ellipsis site (PI; see example 1a in Table 1) and VPE with material between the antecedent and the ellipsis site (RI; see example 1c in Table 2). We used different types of interfering material to explore further the nature of the retrieval process, as discussed earlier. Additionally, different types of interfering material increased the likelihood of detecting different search strategies. That is, a search could in principle be performed over

different types of units, so it is unclear a priori what counts as a relevant constituent in any potential search.

PI was induced by placing before the antecedent (a) a temporal adjunct (e.g., *sometime in the early morning yesterday*); (b) a different cue verb phrase (VP; e.g., *even though Claudia was not particularly angry* in the presence of a *did* cue); (c) a matching cue VP (e.g., *even though Claudia did not write an angry letter* in the presence of a *did* cue); or (d) a clause in which an additional VP overlapped with the retrieval cue, which results in ambiguity (e.g., *Claudia wrote an angry letter, and she also filed a complaint* in the presence of a *did* cue). Under PI, the antecedent (e.g., *filed a complaint*) was adjacent to the VPE (e.g., *Ron did too.*). RI was induced by placing the clauses from the PI conditions at the beginning of the second sentence but after the antecedent. Under RI, the antecedent (e.g., *filed a complaint*) was in the first sentence (e.g., *Claudia filed a complaint*).

We created a matching unacceptable condition, (b/d), by replacing the VPE retrieval cue (e.g., *did too*) with a mismatching cue (e.g., *was too*), which rendered the ellipsis uninterpretable (see Tables 1 and 2). Half of the items used *did* as the acceptable retrieval cue and half used *was*, so that participants could not anticipate stimulus acceptability from the retrieval cue. This also ensured that, to discriminate acceptable from unacceptable sentences, the VPE would have to be processed to the point that the antecedent was retrieved and interpreted.

We included an equal number of acceptable and unacceptable controls, without VPE, such as Examples 5–8 and 13–16 in Tables 1 and 2. These sentences had the same lexical content as Examples 1–4 and 9–12, except that a noun phrase was added to the final

Table 1
Proactive Interference

Examples of proactive interference	
Example item with “did” cue	
VPE	(1a/b) Sometime in the early morning yesterday, Claudia filed a complaint. Ron did too./was ^a too. (2a/b) Even though Claudia was not particularly angry, she filed a complaint. Ron did too./was ^a too. (3a/b) Even though Claudia did not write an angry letter, she filed a complaint. Ron did too./was ^a too. (4a/b) Claudia wrote an angry letter, and she also filed a complaint. Ron did too./was ^a too.
Control	(5a/b) Sometime in the early morning yesterday, Claudia filed a complaint. Ron did the taxes./was ^a the taxes. (6a/b) Even though Claudia was not particularly angry, she filed a complaint. Ron did the taxes./was ^a the taxes. (7a/b) Even though Claudia did not write an angry letter, she filed a complaint. Ron did the taxes./was ^a the taxes. (8a/b) Claudia wrote an angry letter, and she also filed a complaint. Ron did the taxes./was ^a the taxes.
Example item with “was” cue	
VPE	(9a/b) During the morning and into the afternoon, Janice was very angry. David was too./did ^a too. (10a/b) Even though Janice did not write a letter, she was very angry. David was too./did ^a too. (11a/b) Even though Janice was not particularly frustrated, she was very angry. David was too./did ^a too. (12a/b) Janice was exhausted from work, and she also was very angry. David was too./did ^a too.
Control	(13a/b) During the morning and into the afternoon, Janice was very angry. David was annoyed./did ^a annoyed. (14a/b) Even though Janice did not write a letter, she was very angry. David was annoyed./did ^a annoyed. (15a/b) Even though Janice was not particularly frustrated, she was very angry. David was annoyed./did ^a annoyed. (16a/b) Janice was exhausted from work, and she also was very angry. David was annoyed./did ^a annoyed.

Note. VPE = verb-phrase ellipsis.

^a Unacceptable.

Table 2
Retroactive Interference

Examples of retroactive interference	
Example item with “did” cue	
VPE	(1c/d) Claudia filed a complaint. Sometime in the early morning yesterday, Ron did too./was ^a too. (2c/d) Claudia filed a complaint. Even though he was not particularly angry, Ron did too. /was ^a too. (3c/d) Claudia filed a complaint. Even though he did not write an angry letter, Ron did too. /was ^a too. (4c/d) Claudia filed a complaint and she also wrote an angry letter. Ron did too./was ^a too.
Control	(5c/d) Claudia filed a complaint. Sometime in the early morning yesterday, Ron did the taxes./was ^a the taxes. (6c/d) Claudia filed a complaint. Even though he was not particularly angry, Ron did the taxes./was ^a the taxes. (7c/d) Claudia filed a complaint. Even though he did not write an angry letter, Ron did the taxes./was ^a the taxes. (8c/d) Claudia filed a complaint and she also wrote an angry letter. Ron did the taxes./was ^a the taxes.
Example item with “was” cue	
VPE	(9c/d) Janice was very angry. During the morning and into the afternoon, David was too./did ^a too. (10c/d) Janice was very angry. Even though he did not write a letter, David was too./did ^a too. (11c/d) Janice was very angry. Even though he was not particularly frustrated, David was too./did ^a too. (12c/d) Janice was very angry and she was exhausted because of work. David was too./did ^a too.
Control	(13c/d) Janice was very angry. During the morning and into the afternoon, David was annoyed./did ^a annoyed. (14c/d) Janice was very angry. Even though he did not write a letter, David was annoyed./did ^a annoyed. (15c/d) Janice was very angry. Even though he was not particularly frustrated, David was annoyed./did ^a annoyed. (16c/d) Janice was very angry and she was exhausted because of work. David was annoyed./did ^a annoyed.

Note. VPE = verb-phrase ellipsis.

^a Unacceptable.

clause to block elliptical interpretation. They were included to prevent anticipation of VPE from the initial sentence form.

Participants read 64 experimental sentences per session, counter-balanced within and across sessions, and saw every item in both conditions at different points. Critical trials, including unelided controls, constituted 22% of each session and were presented randomly among the remaining stimuli, none of which was elided. The fillers were multiclausal sentences, with equal numbers of acceptable and unacceptable versions: *The doctor who had wondered if the nurse was on break heard that the patient cried/creased (creased is unacceptable).*

Norming

We collected plausibility ratings from 20 other (paid) participants. Participants rated stimuli on a scale from 1 (*highly implausible*) to 7 (*highly plausible*), with 4 being neutral. The stimuli were pseudorandomly interleaved with fillers, such that no item appeared more than twice and no condition appeared more than three times, consecutively. Because we aimed to encourage the full use of the rating scale, 25% of the filler items were highly implausible. Elided conditions were rated as plausible for both RI and PI conditions, RI by participants $M = 5.51$ ($SD = 0.70$), RI by items $M = 5.51$ ($SD = 0.47$), PI by participants $M = 5.48$ ($SD = 0.69$), PI by items $M = 5.54$ ($SD = 0.46$). No significant differences were found. This indicates that potential differences between RI and PI conditions in accuracy would be due to the interference manipulation but not to differences in plausibility.

Procedure

Participants read sentences phrase-by-phrase in a controlled presentation manner, 335 ms per word in the phrase. Participants were trained to respond to a signal (50 ms, 1,000 Hz tone) presented 14 times after the onset of a critical expression, here VPE. The first response signal occurred 300 ms before the onset of the VPE, and thus participants responded before processing of the critical expression began. The subsequent sampled times (0–4,900 ms) enabled us to measure how the interpretation of the VPE unfolded over time. They simultaneously pressed both *yes* and *no* keys as an initial (undecided) response and then switched to one key when information regarding the acceptability of the sentence became available. Participants were encouraged to modulate their responses if their judgment changed during the trial. Stimulus presentation and response collection were carried out with software with millisecond timing. Between-trial intervals were participant controlled, and there were two mandatory breaks each session.

Data Analysis

Comprehension accuracy was calculated using a standard d' , $d' = z(\text{hits}) - z(\text{false alarms})$, where a “hit” was an “acceptable” response to an acceptable sentence and a “false alarm” was an “acceptable” response to an unacceptable sentence. Differences between conditions were quantified with hierarchically nested fits of the exponential (Equation 1) to each participant’s data. Models ranged from a null model (all conditions fit with $1\lambda - 1\beta - 1\delta$) to a

fully saturated model (unique parameters for each condition), with the quality of fits being assessed by an adjusted multiple correlation-squared (adjusted R^2) statistic, crucially by evaluating the consistency of parameter estimates across participants, and by inferential significance tests.

Results

Figure 2 presents the average (across participants) d' as a function of processing time, along with the best-fitting exponential model described below. Inspection of Figure 2 suggests that VPE under RI was less accurately processed than VPE under PI. No systematic differences were found between types of interfering material (see Table 3 for means). As such, we analyzed all PI and RI conditions together, respectively, to test for effects of PI versus RI. We averaged the d' values for each participant (and by each item for an item analysis) in each condition from 3.5 to 4.9 s postinitial response cue to derive an empirical estimate of asymptotic accuracy. Responses to elided sentences with PI were on average 0.37 d' units higher in accuracy than responses to elided sentences with RI (95% CI = 0.19, 0.54 d' units). A paired t test¹ showed that this difference was significant, $F_1(1, 11) = 21.34, p < .001$, and $F_2(1, 35) = 8.18, p < .01$, $\min F'(1, 46) = 5.91, p < .05$. Accuracy for RI unelided sentences was 0.06 d' units lower than accuracy for PI unelided sentences (95% CI = -0.05, 0.17 d' units), which was not significant.

We performed a repeated measures analysis of variance (ANOVA) on elided and control conditions. Though the controls do not form a true baseline of the elided conditions, this analysis showed that the location of interfering material affected elided sentences more than in control versions. There was a main effect of ellipsis and of retrieval context, as well as an interaction (see Table 4). Pairwise comparisons confirmed that the source of the interaction was a difference between the elided PI and RI conditions (3.63 vs. 3.26, 95% CI = 0.19, 0.54). Crucially, there was no difference between PI and RI controls (3.84 vs. 3.78, 95% CI = -0.05, 0.17), indicating that retrieval context negatively affected the probability of successful VPE interpretation.

Competitive fits also yielded clear evidence that retrieval context modulated asymptotic performance: Models that did not allocate separate asymptotes for PI versus RI ellipsis conditions pro-

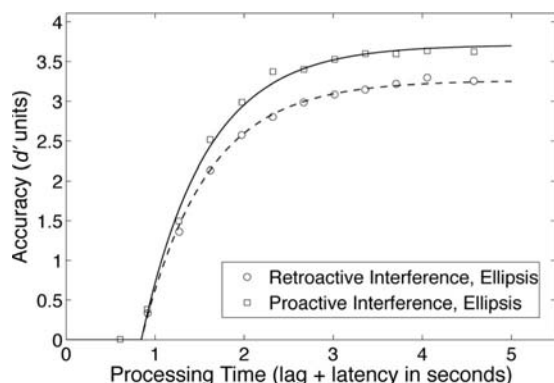


Figure 2. Average d' accuracy (symbols) as a function of processing time (lag of the interruption response-signal plus response latency) for proactive interference (PI) and retroactive interference (RI) elided conditions. Smooth curves show the best fitting exponential fit.

duced poor fits and left systematic residuals. In fits of the average data, allocating separate asymptotes to each ellipsis condition increased adjusted R^2 from .973 observed with a null $1\lambda-1\beta-1\delta$ model to .997. This $2\lambda-1\beta-1\delta$ model improved the quality of the fits of the individual participants' data, systematically increasing adjusted R^2 over what was observed with a $1\lambda-1\beta-1\delta$ model (ranging from .973-.993, compared with .832-.985). In the average data, the asymptote for VPE under PI was estimated to be 3.74, whereas the estimate for VPE under RI was 3.29. Across participants, the average difference in asymptotic (λ) estimates was 0.45 d' units (95% CI = 0.26, 0.63 d' units), which a paired t test showed was significant, $F(1, 11) = 28.72, p < .001$.

The differences in asymptote indicate that VPE under RI was less accurately processed than VPE under PI, or that quality of the retrieved information was poorer for antecedents under RI, leading to less acceptable interpretation. Crucially, however, if the PI or RI contexts affected processing speed, then it should have engendered differences in either rate (β) or intercept (δ). It is important that allocating separate β or δ parameters to PI and RI conditions did not improve adjusted R^2 . In fits of the average, a $2\lambda-2\beta-1\delta$ model resulted in an adjusted R^2 of .998 and a $2\lambda-1\beta-2\delta$ model resulted in adjusted R^2 of .997, compared with the .997 value observed with the simpler $2\lambda-1\beta-1\delta$ model. It is also important that there were no consistent trends across participants in either rate or intercept parameters when they were allowed to vary, and tests on the parameter estimates were not significant.² Therefore, there was no evidence suggesting that PI or RI affected processing speed.

For completeness, we also compared the functions for controls without VPE. As Figure 3 suggests, no differences were found. Consequently, the best fit for these functions was a simple $1\lambda-1\beta-1\delta$ model, adjusted $R^2 = .993$. All t tests on parameter estimates for models that varied one of the SAT parameters were not

¹ For $\min F'$ values, we computed the F statistic as the square of the t statistic.

² Rate differences can sometimes trade off with asymptotic differences in model fits. As an illustration, consider a case where the true state of affairs (or true model fit) consists of two functions rising to the same asymptote at different rates, where Condition X rises faster than Condition Y. The rise part of the function, or the portion of the function based on the earliest data points, will be higher for Condition X than for Condition Y and would ideally be fit with a faster rate parameter for the former condition. However, the higher valued early points on the function fit to Condition X could be fit or captured by the slower rate parameter of Condition Y, if the asymptote of Condition X were set to a higher value. This is because the rate parameter determines what proportion of the asymptote is reached at any point in time. The least squares procedure uses an omnibus statistic, minimizing the deviation of predicted from observed points across the entire function. If there are too few points on the asymptote portion of the function, or if there is some nonmonotonicity in the data, a better fit by an omnibus criterion might result from (in this illustration, inappropriately) varying the asymptotic parameter rather than the rate parameter. Could the uniform dynamics observed here have resulted from such a tradeoff? Inspection of the best fitting $2\lambda-1\beta-1\delta$ model does not show any clear evidence of overestimated asymptotes that are characteristic of this type of tradeoff. Nonetheless, to rigorously test for this possibility, we also performed the series of fits with a fixed asymptote, where the highest possible value for the λ parameter was fixed to the final empirical d' value for each condition. This strong test for dynamics differences did not change the pattern of results reported above; no reliable differences in either rate or intercept were found between conditions.

Table 3
Accuracy by Type of Interfering Material

Empirical d' by condition for elided conditions	M (SE) in d' units
Under RI	
Temporal adjunct	3.34 (0.16)
Different cue	3.34 (0.14)
Same cue	3.22 (0.16)
Ambiguous cue	3.13 (0.15)
Mean across conditions	3.26 (0.14)
Under PI	
Temporal adjunct	3.67 (0.08)
Different cue	3.63 (0.14)
Same cue	3.67 (0.11)
Ambiguous cue	3.54 (0.16)
Mean across conditions	3.63 (0.11)

Note. RI = retroactive interference; PI = proactive interference.

significant, suggesting that the interference source effect evident in VPE is related to the retrieval of the antecedent, not to general differences between the PI and RI sentence forms.

Discussion

Successful comprehension often requires dependency formation between nonadjacent constituents, of which an essential component is access to memory representations outside of the current focus of attention. In this experiment, we tested whether forward or backward search mediates VPE resolution by examining comprehension speed as a function of interference in the retrieval context. We equated the overall contents of memory, so that any observed differences cannot be attributed to the presence of additional material alone (as in a standard distance manipulation).

If forward search mediates access to the antecedent, then PI induced by additional material before the antecedent should have slowed retrieval speed, which in turn should have slowed processing. Conversely (contrary to the results in Martin & McElree, 2008), if backward search is used, then RI induced by materials between the antecedent and ellipsis should have slowed processing. We found that there were no differences in processing speed between PI and RI contexts, which provides evidence against a (forward or backward) search. This pattern held across types of interfering material, suggesting that the results are robust across different potential ways in which the search units can be defined.

The observed processing dynamics are predicted by a direct-access process, which enables representations of differing quality to be recovered with similar dynamics (McElree, 2006). Direct-

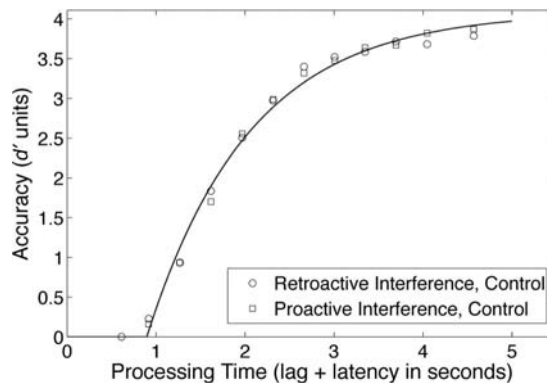


Figure 3. Average d' accuracy (symbols) as a function of processing time (lag of the interruption response-signal plus response latency) for proactive interference and retroactive interference control conditions. Smooth curves show the best fitting exponential fit.

access retrieval implies a very high degree of content addressability of the representations formed during comprehension, as it suggests that the retrieval process delivers the target via its content alone (Martin & McElree, 2008). The present results are consistent with other findings indicating content-addressable representations are formed during sentence comprehension (Martin & McElree, 2008; McElree, 2000, 2006; McElree et al., 2003). Previous findings were derived solely from manipulations of RI, so those studies do not address the possibility of forward search. Our data indicate that neither forward nor backward search, even if defined over various possible constituent types, provides an adequate account of ellipsis resolution.

Furthermore, our data indicate that although additional information in memory did not affect retrieval speed, it does affect asymptotic accuracy, with lower levels of accuracy for RI than PI conditions. There are several possible explanations for this effect, all of which are consistent with content addressability and direct access. Processing material after the antecedent could reduce the quality of antecedent encoding, or, in the limit, displace it from memory altogether. However, most basic memory research indicates that interference primarily affects retrieval (e.g., Anderson & Neely, 1996; Crowder, 1976); specifically, it induces cue overload such that extant retrieval cues are insufficient for successful elicitation of the target (Nairne, 2002; Öztekin & McElree, 2007; Watkins & Watkins, 1975). In line with these general approaches and with studies showing detrimental effects of similarity-based interference in comprehension (e.g., Gordon, Hendrick, & Levine, 2002; Van Dyke, 2007; Van Dyke & Lewis, 2003; Van Dyke &

Table 4
Retrieval Context \times Ellipsis Analysis of Variance on Empirical d'

Pair	M difference (SE)	Source	$F_1(df)$	F_1	$F_2(df)$	F_2	$minF'(df)$	$minF'$
PI vs. RI, ellipsis	.37 (.08)	Ellipsis	1, 11	9.48**	1, 35	16.55**	1, 25	6.02*
		Retrieval context	1, 11	16.82**	1, 35	8.15**	1, 44	5.49*
PI vs. RI, control	.06 (.05)	Retrieval Context \times Ellipsis	1, 11	14.12**	1, 35	4.05*	1, 46	3.14 [†]

Note. PI = proactive interference; RI = retroactive interference.
[†] $p < .09$. * $p < .05$. ** $p < .01$.

McElree, 2006), we assume that the observed asymptotic differences are due to differences at retrieval. Specifically, it is likely that RI had a more detrimental effect on comprehension because recently processed material is more salient at retrieval and thus more likely to inappropriately intrude upon retrieval of the target.

There are two ways in which the PI context benefits from recency. First, the antecedent itself was more recent under PI than under RI, as it was processed immediately before the ellipsis. Thus, the memory strength of the antecedent at retrieval under PI was likely stronger than under RI. Second, the interfering material was more recent under RI than under PI, which would increase cue diagnosticity in the latter. Both factors are likely to have conspired to increase the resonance of retrieval cues with the memory representation of the antecedent under PI relative to RI.

Beyond the sizable asymptotic differences between the PI and RI conditions, we did not observe differential effects of type of interfering material within the RI and PI conditions. This might appear *prima facie* inconsistent with cue-dependent direct access. However, with one exception, none of the interpolated constituents used in the experiment were fully acceptable antecedents for the ellipsis, and it possible that the semantic/pragmatic and morphosyntactic cues at the ellipsis site were specific enough to avoid strongly and differentially eliciting any of the competing structures of the type explored here. For example, the temporal adjunct (e.g., *sometime in the early morning yesterday*) in Examples 1 and 9 in Tables 1 and 2 lacks VP status, and both of the VP structures (e.g., *even though Claudia was not particularly angry* and *even though Claudia did not write an angry letter*) in Examples 2 and 10, and 3 and 11, respectively, in Tables 1 and 2 involve negation, which mismatches the affirmative properties of the antecedent. The only condition with full overlap with the retrieval cues was the ambiguous case (Examples 4 and 12 in Tables 1 and 2). This type of constituent may not have delayed processing or resulted in retrieval of the incorrect target because assigning either constituent to the VPE, or even the conjunction of the constituents, would result in an acceptable elided interpretation. Thus, rather than being inconsistent with a content-addressable process, the current finding may indicate that, beyond the general effects of recency that engendered differences between the RI and PI conditions, retrieval cues at the ellipsis site in the current study were distinctive enough to point to the correct antecedent in memory, without substantial competition from other seemingly similar constituents in memory.

Clearly, additional work is needed to identify the cues used to recover the antecedents of VPEs and, thus, what types of constituents are likely to engender interference. Investigations of nonadjacent verb-argument dependencies have shown that interference is modulated by the overlap between the semantic and syntactic features of constituents in memory and the to-be-retrieved constituent (e.g., Van Dyke, 2007). To the degree that VPEs are processed with the same type of direct-access operation, we would expect that a thorough investigation of competing constituents would likewise uncover differential interference effects in VPE processing. We note, however, that the relationship between the ellipsis site and the antecedent in memory has been argued to be more anaphoric in nature than verb-argument dependencies (Hardt, 1999). If correct, future research may discover that VPEs are susceptible to different kinds of interference than other types of nonadjacent dependencies.

In summary, our results indicate that the time-course profile for retrieval operations used in ellipsis resolution is identical to what

has been observed in the processing of other nonadjacent dependencies: Although the recency of the antecedent affects the likelihood of successful interpretation, lowering asymptotic accuracy, it does not affect the speed with which an antecedent representation is accessed. The unique contribution of the reported study is that the time-course profiles for ellipsis resolution are incompatible with classes of both backward and forward search mechanisms but are consistent with previous findings suggesting that linguistic representations are content addressable and accessed with a cue-dependent direct-access operation.

References

- Anderson, M. C., & Neely, J. H. (1996). Interference and inhibition in memory retrieval. In E. L. Bjork & R. A. Bjork (Eds.), *Handbook of perception and memory: Vol. 10. Memory* (pp. 237–313). San Diego, CA: Academic Press.
- Clark, S. E., & Gronlund, S. D. (1996). Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin & Review*, 3, 37–60.
- Crowder, R. G. (1976). *Principles of learning and memory*. Hillsdale, NJ: Erlbaum.
- Doshier, B. A. (1979). Empirical approaches to information processing: Speed-accuracy tradeoff functions or reaction time. *Acta Psychologica*, 43, 347–359.
- Gernsbacher, M. A., & Hargreaves, D. (1988). Accessing sentence participants: The advantage of first mention. *Journal of Memory and Language*, 27, 699–717.
- Gordon, P. C., Hendrick, R., & Levine, W. H. (2002). Memory-load interference in syntactic processing. *Psychological Science*, 13, 425–430.
- Gronlund, S. D., Edwards, M. B., & Ohrt, D. D. (1997). Comparison of the retrieval of item versus spatial position information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 1261–1274.
- Hardt, D. (1999). Dynamic interpretation of verb phrase ellipsis. *Linguistics and Philosophy*, 22, 187–221.
- Martin, A. E., & McElree, B. (2008). A content-addressable pointer underlies comprehension of verb-phrase ellipsis. *Journal of Memory and Language*, 58, 879–906.
- McElree, B. (1993). The locus of lexical preference effects in sentence comprehension: A time-course analysis. *Journal of Memory and Language*, 32, 536–571.
- McElree, B. (2000). Sentence comprehension is mediated by content-addressable memory structures. *Journal of Psycholinguistic Research*, 29, 111–123.
- McElree, B. (2001). Working memory and focal attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 817–835.
- McElree, B. (2006). Accessing recent events. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 46). San Diego, CA: Academic Press.
- McElree, B., & Doshier, B. A. (1989). Serial position and set size in short-term memory: Time course of recognition. *Journal of Experimental Psychology: General*, 18, 346–373.
- McElree, B., & Doshier, B. A. (1993). Serial retrieval processes in the recovery of order information. *Journal of Experimental Psychology: General*, 122, 291–315.
- McElree, B., Foraker, S., & Dyer, L. (2003). Memory structures that subserve sentence comprehension. *Journal of Memory and Language*, 48, 67–91.
- Murphy, G. L. (1985). Processes of understanding anaphora. *Journal of Memory and Language*, 24, 290–303.
- Nairne, J. S. (2002). The myth of the encoding-retrieval match. *Memory*, 10, 389–395.

- Öztekın, I., & McElree, B. (2007). Retrieval dynamics of proactive interference: PI slows retrieval by eliminating fast assessments of familiarity. *Journal of Memory and Language*, *57*, 126–149.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.
- Reed, A. V. (1973). Speed-accuracy trade-off in recognition memory. *Science*, *181*, 574–576.
- Thomas, J. G., Milner, H. R., & Haberlandt, K. F. (2003). Forward and backward recall: Different response time patterns, same retrieval order. *Psychological Science*, *14*, 169–174.
- Townsend, J. T., & Ashby, F. G. (1983). *The stochastic modeling of elementary psychological processes*. New York: Cambridge University Press.
- Van Dyke, J. A. (2007). Interference effects from grammatically unavailable constituents during sentence processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 407–430.
- Van Dyke, J. A., & Lewis, R. L. (2003). Distinguishing effects of structure and decay on attachment and repair: A cue-based parsing account of recovery from misanalyzed ambiguities. *Journal of Memory and Language*, *49*, 285–316.
- Van Dyke, J. A., & McElree, B. (2006). Retrieval interference in sentence comprehension. *Journal of Memory and Language*, *55*, 157–166.
- Watkins, O. C., & Watkins, M. J. (1975). Build-up of proactive inhibition as a cue overload effect. *Journal of Experimental Psychology: Human Learning and Memory*, *104*, 442–452.
- Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica*, *41*, 67–85.

Received July 22, 2008

Revision received March 17, 2009

Accepted March 25, 2009 ■