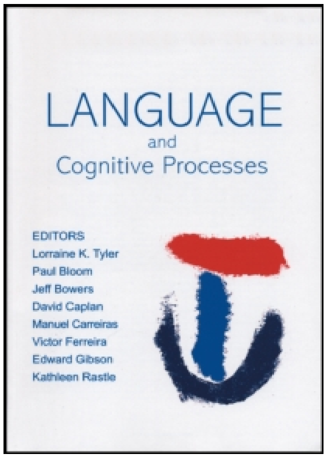

Language and Cognitive Processes
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713683153

Effects of NP type in reading cleft sentences in English

Tessa Warren [a]; Edward Gibson [b]
[a] University of Pittsburgh, Pittsburgh, PA, USA
[b] Massachusetts Institute of Technology, Cambridge, MA, USA

Online Publication Date: 01 December 2005
To cite this Article: Warren, Tessa and Gibson, Edward (2005) 'Effects of NP type in reading cleft sentences in English', Language and Cognitive Processes, 20:6, 751 - 767
To link to this article: DOI: 10.1080/01690960500051055
URL: http://dx.doi.org/10.1080/01690960500051055

LANGUAGE AND COGNITIVE PROCESSES
2005, 20 (6), 751–767

Psychology Press
Taylor & Francis Group

Effects of NP type in reading cleft sentences in English

Tessa Warren
University of Pittsburgh, Pittsburgh, PA, USA

Edward Gibson
Massachusetts Institute of Technology, Cambridge, MA, USA

This paper investigates factors which contribute to the complexity of English sentences with long-distance dependencies. Two hypotheses were compared: (1) increased referential processing between the endpoints of a dependency increases processing difficulty at the completion of the dependency (Gibson, 1998, 2000; Warren & Gibson, 2002) and (2) increased similarity between NPs awaiting role-assignment increases memory interference during retrieval (Gordon, Hendrick, & Johnson, 2001). Self-paced word-by-word moving-window reading times were gathered over object-extracted cleft sentences in which two NPs were varied among definite descriptions, first names, and pronouns. Reading times at the verb supported both hypotheses. As the referential hypothesis predicted, reading times were faster when the intervening subject NP had a more referentially accessible type. Consistent with the similarity hypothesis, reading times were slow when both NPs were names or descriptions. Later comprehension measures showed strong effects of similarity-based interference, but did not show effects of referential processing load.

An aim of human sentence processing research is to discover what kinds of information people use and what computations they perform in the moment-by-moment comprehension of a sentence. Research has shown that during comprehension, people use information from many sources,

Correspondence should be addressed to either author: Tessa Warren, 607 LRDC, 3939 O'Hara St., Pittsburgh, PA 15260, USA. Email: tessa@pitt.edu; Edward Gibson, NE20-459, MIT, Cambridge, MA 02139, USA. Email: egibson@mit.edu

This work was undertaken while the first author was supported as a post-doctoral fellow on NIMH training grant MH16745 awarded to the Department of Psychology at the University of Massachusetts. The authors would like to thank Mariko Jameson for help writing the stimuli, and Christopher Hirsch and the audiences of AMLaP 2003 and NELS 2003 for helpful comments on this work.

http://www.psypress.com.lcp DOI: 10.1080/01690960500051055

including the lexicon, syntax, the discourse context and world knowledge (for reviews, see Gibson & Pearlmutter, 1998; Tanenhaus & Trueswell, 1995). This information is put to many uses, including aiding in the computation of the inter-word dependencies that determine syntactic agreement and semantic interpretation. The current paper investigates factors influencing the difficulty of computing these dependencies during comprehension.

A sentence's dependency structure has been shown to be a strong determinant of its processing difficulty (Gibson, 1998, 2000; Grodner & Gibson, 2005). An example supporting this is the contrast between object- and subject-extracted relative clauses (RCs) in Subject-Verb-Object languages like English, as in (1):

(1) a. Object-extraction: The reporter [that the senator attacked] disliked the editor.
b. Subject-extraction: The reporter [that attacked the senator] disliked the editor.

RCs are noun-modifying clauses, exemplified by the bracketed clauses in (1a) and (1b). In (1a) there is a dependency between *that* and the object position of the verb *attacked*, because the relative pronoun is object-extracted—it is interpreted as the object of *attacked*. Correspondingly, the dependency between *that* and *attacked* in (1b) is a subject-extraction because *that* is interpreted as the subject of the verb *attacked*. Object-extractions cause increased processing difficulty as measured by phoneme-monitoring, on-line lexical-decision latency, reading times, and response accuracy to probe questions (Hakes, Evans, & Brannon, 1976; Wanner & Maratsos, 1978; Holmes & O'Regan, 1981; Ford, 1983; Waters, Caplan, & Hildebrandt, 1987; King & Just, 1991). This effect is not driven by lexical frequencies, because the sentences have the same words, nor real-world plausibility, because experiments have controlled for plausibility and still found the effect. Most experiments test these sentences in a null context, ruling out the possibility that differences in discourse context drive the effect. The effect must therefore be related to the different word orders in the sentences. There are multiple theories of why the word order in object-extracted RCs results in more processing difficulty than the word order in subject-extracted RCs, including ones that attribute the additional processing load to syntactic storage (Chomsky & Miller, 1963; Gibson, 1998), structural frequency distributions (Bever, 1970; MacDonald & Christiansen, 2002) or perspective shift (MacWhinney & Pleh, 1988). In the current paper, we evaluate theories attributing these processing complexity differences to the differences in dependency structure between object and subject extractions. This paper investigates the relation between dependencies and processing difficulty, by examining how people process variants of sentences like (1a).

A number of theories attribute the difference in processing difficulty between object- and subject-extractions to factors affecting dependency calculation and integration (Gibson, 1998, 2000; Gordon, Hendrick, & Johnson, 2001, 2004; Van Dyke & Lewis, 2003). Integration is the process of creating dependencies by linking a new word into a sentence representation, which requires retrieval of the previously processed material upon which the new word is dependent. In ambiguity resolution, dependencies spanning shorter distances are preferred over those spanning longer distances (Kimball, 1973; Frazier, 1987; Gibson, 1991; Pearlmutter & Gibson, 2001). Additionally, reading times on unambiguous structures suggest that the processing difficulty caused by an integration is correlated with the distance it spans (Gibson, 1998, 2000; Grodner & Gibson, 2005). These findings suggest that the reason object-extracted RCs are more difficult may be that they involve longer dependencies than subject-extracted RCs. Whereas some theories assume that the relevant integrations are syntactic (e.g., Gibson, 1998, 2000; Van Dyke & Lewis, 2003), others consider them to be semantic (Gordon et al., 2001, p. 1417). However, no empirical results here or in previous work depend on either initial assumption. In the object-extracted RC in (1a), connecting the verb *attacked* into the sentence involves integrating the relative pronoun *that* as the object of *attacked*, across the noun phrase (NP) *the senator*. In contrast, all integrations in the subject-extracted RC in (1b) are between consecutive words.

Although dependency length seems to predict processing difficulty, there is no consensus yet about how to quantify length. Gibson (1998, 2000) suggests that one possible cause for integration difficulty is that the activation level of earlier-processed material may decay as later information is processed. Processing each new word may cause the representation of earlier attachment sites to decay, supporting a word-based length metric (cf. Hawkins, 1994). Gibson (1998, 2000) and Warren and Gibson (2002) pursue another potential length metric; one that is discourse referent-based. Introducing new referents imposes a processing cost (Murphy, 1984; Haviland & Clark, 1974). Warren and Gibson hypothesise that when more working memory resources are directed to referential processing, fewer resources are available for maintaining syntactic representations and they decay. The prediction then, is that costs resulting from the referential processing of an NP intervening between the endpoints of a dependency may be detected more than once, first at the NP and then again at the conclusion of the dependency. Consistent with this hypothesis, integrations that cross indexical pronouns—first- and second-person pronouns such as ‘‘I’’, ‘‘you’’ and ‘‘we’’, whose referents are anchored in the context and thus require less referential processing—cause less processing difficulty than integrations crossing new referents (Warren & Gibson, 2002). For example, replacing the embedded subject ''the senator'' in sentence (1a) with the pronoun ''you'' leads to faster reading times at the embedded and main verbs, where the filler-gap and subject-verb integrations are completed (Warren & Gibson, 2002; Gordon et al., 2001).

Warren and Gibson (2002) extended this finding across multiple types of NPs. In sentences like (2), which were controlled for plausibility, the verbs were read faster when the RC subject position was occupied by a NP type closer to the given end of a Givenness Hierarchy of nominal reference (Ariel, 1990; Gundel, Hedberg, & Zacharski, 1993).

(2) The consultant who (we/Donald Trump/the chairman/a chairman) called advised wealthy companies.

Givenness hierarchies correlate NP types with the accessibility of their referents. Pronouns, such as ''we'' in (2), always refer to highly accessible referents, while first or famous names [''Donald Trump'' in (2)] are used for slightly less available referents and definite descriptions [''the chairman'' in (2)] for even less accessible referents. Indefinite descriptions [''a chairman'' in (2)] introduce new referents. Warren and Gibson hypothesised that for NPs with referents in the discourse, their type may predict the ease of accessing their referents. For NPs without referents in the discourse, the effort spent attempting to access a referent before quitting and instantiating a new one may be less for NP types whose referents must be highly activated than for NP types whose referents are usually less activated. Thus NP processing cost is expected to be least for pronouns, higher for names, and higher still for definite descriptions, both in null context and in natural text. We will refer to this hypothesis as the referent accessibility integration cost hypothesis. Warren and Gibson (2002) predicted that this NP processing cost would affect the difficulty of integrating across the manipulated NP. They found the expected trend: reading times on the main verb were fastest in the pronoun condition, slower for names, slower for definite descriptions, and slowest for indefinite descriptions.[1]

Another potential source of integration difficulty is interference among similar elements in memory during the retrieval of the integration

[1] The referent-accessibility account actually predicts less difficulty for indefinites than definites, as indefinites do not require a search for a referent or accommodation. Warren and Gibson (2002) provided two possible explanations for the unexpected difficulty: 1) Indefinite NPs are unusual in restrictive RCs which usually contain background information, 2) Building a representation for a new referent introduced with an indefinite may be harder than for a definite (Webber, 1979).

attachment site (Gordon et al., 2001; Lewis, 1996; Van Dyke & Lewis, 2003). According to the similarity based integration cost hypothesis, integrating a word into a sentence representation is predicted to be more difficult when there are multiple words or sub-structures in memory similar to the ones that must be retrieved and integrated with the new word.

The similarity based integration cost hypothesis provides an alternative explanation for the easier processing of object-extracted RCs with pronominal subjects (Gordon et al., 2001). Object-extracted RCs with the same type of NP in the head and subject positions are difficult to process because similarity based interference hampers retrieval of the correct NPs at the verbs. However, the appropriate characterisation of NP type has not been fully determined. Gordon, Hendrick, and Johnson (2004) began to address this question by comparing NPs with different semantic properties, and found that while first names, second person pronouns and the quantified NP ‘‘everyone’’ did not interfere with definite descriptions, every type of NP with a common noun they tested did. Therefore, it seems that the relevant metric for determining the similarity of NP types does not weight semantic factors such as definiteness or genericness, but instead includes factors such as whether the NP includes a common noun.

The similarity based hypothesis can account for Warren and Gibson’s finding that for RCs modifying a definite description, reading times on the verbs were faster when the RC subject was a pronoun or name rather than a description. This is because there is less interference between a pronoun or name and description than between two descriptions.

Gordon et al. (2001) presented four self-paced reading experiments, three of which were consistent with both the similarity based and referential accessibility integration cost hypotheses. They found that object-extracted RCs headed by a description were easier to process if the extraction crossed a pronoun or name rather than another description [as in (2)]. While Gordon et al. did not directly compare names to pronouns, they ran items in separate experiments that differed only in whether the non-case-marked pronoun *you* or a three-letter name interrupted the extraction. The similarity based hypothesis makes no prediction about differences between names and pronouns in these cases, but the referential accessibility hypothesis predicts less difficulty in the pronoun than name conditions. In fact, reading times on the critical word were 190 ms faster in the pronoun condition than the name condition, consistent with the referential accessibility hypothesis.

However, the results of Gordon et al.’s Experiment 4 were consistent only with the similarity based hypothesis. This experiment involved clefted materials as in (3):

(3) a. Subject-extractions: It was (the barber/John) that saw (*the lawyer/Bill*) in the parking lot.
b. Object-extractions: It was (the barber/John) that (the lawyer/ Bill) *saw* in the parking lot.

The experiment was a 2 × 2 × 2 design, crossing extraction type (subject- vs. object-extraction), clefted NP type (definite description, name), and embedded NP type (definite description, name). Both theories predicted object-extractions would be more difficult to process than subject-extractions because they involve longer dependences, and subject-extractions would vary less because of their shorter integrations. As predicted, reading times at the critical word [italicised in (3)] were slower for object-extracted clefts than subject-extracted clefts and similar among the subject-extracted conditions. As the similarity based hypothesis predicted, within the object-extracted conditions reading times on the critical word were slower for the matching NP conditions (description-description and name-name) than the non-matching NP conditions (description-name and name-description). Furthermore, the two matching conditions did not differ, nor did the two non-matching conditions. Whereas the similarity based hypothesis predicts exactly this pattern, the referential accessibility hypothesis predicts slower reading times on the critical word in the object-extracted conditions with embedded descriptions, because long dependencies crossing less given NPs are predicted to be more difficult to process. No such effect was observed.

Although the results from Experiment 4 of Gordon et al. (2001) favoured the similarity based hypothesis, Gordon et al. did not test pronominals, perhaps the most critical comparison for evaluating the referential accessibility hypothesis in null context sentences like these where first names and descriptions have no referents. Indexical pronouns map onto default roles included in every communicative event, even when the particular individuals they refer to cannot be identified. Because indexical pronouns can be assigned the generic referents of communicator and communicatee, they provide a test case of an NP with an accessible referent. The current experiment was designed to (1) include the test case of indexical pronouns; and (2) investigate the time course of the effects, because Gordon et al. only reported reading times at the critical word.

EXPERIMENT

This experiment tested nine versions of object-extracted cleft sentences. Subject-extracted cleft sentences were not tested, as they do not differentiate between the similarity based and referential accessibility hypotheses. The experiment had a 3 × 3 design, crossing three types of NP (definite descriptions, names, and pronouns) in the clefted (NP1) and embedded (NP2) positions. An example item is provided in (4):

(4) It was (the lawyer/Patricia/you) who (the businessman/Dan/we) avoided at the party.

The similarity based hypothesis predicts a pattern similar to the results of Experiment 4 in Gordon et al. (2001): Conditions with matching NPs (i.e. pronoun-pronoun, name-name and description-description) should show the longest reading times at the verb.

In contrast, the referential accessibility hypothesis predicts that reading times on the verb should vary with the position of NP2 on the Givenness Hierarchy, because NP2 intervenes between the endpoints of the verb-object integration, while NP1 does not. The referential accessibility hypothesis predicts a main effect of NP2 type, with the longest reading times when NP2 is a description, faster when it is a name, and fastest when it is a pronoun. This hypothesis makes no predictions based on the similarity of the two NP types.

Method

Participants. Forty-two members of the MIT community were paid $10 each to participate in this experiment.

Materials. There were 36 items having the form of (4), in the design described above. Eighteen were modified versions of items from Gordon et al.'s (2001) Experiment 4. An additional 18 items were constructed in the same style. Experimental items are included in Appendix A.

Only first- or second-person pronouns were used. The pronoun in the NP1 position was always in accusative case. An equal number of items used ''us'', ''you'', and ''me'' as NP1. The pronoun in the NP2 position was always in nominative case. An equal number of items used ''I'', ''we'' and ''you'' as NP2. The items were further constrained so that none of the pronoun-pronoun items contained both ''us'' and ''we'', both ''you'' and ''you'', or both ''me'' and ''I''. The definite descriptions used in the experimental items were standard occupations or roles for humans, for example: ''customer'', ''doctor'', ''homeowner'', and ''busboy''.

Every sentence in the experiment was followed by a true-false question. Most questions about experimental items tested knowledge of the relation between the NPs and verb, but some tested knowledge of the adjunct phrase that completed the sentence. For example, one condition of (4) was followed by the question: ''Was the businessman avoided at the bar?'' while another condition was followed by ''Did you avoid Dan?''

The 36 items from this experiment were combined with 24 items from each of two unrelated experiments and 36 filler sentences, for a total of 120 sentences. The nine conditions from this experiment were counterbalanced across lists. Every subject saw one version of each item and four versions of each condition across the experiment. The order of sentence presentation was pseudo-randomised for each participant, with the constraint that items from the same experiment not appear consecutively.

Procedure. The task was self-paced, word-by-word reading using a moving window display (Just, Carpenter, & Woolley, 1982).[2] Sentences were initially displayed with dashes replacing all letters and punctuation. Participants pressed the spacebar to reveal each word of the sentence, and as they did so the preceding word reverted to dashes. Reading time was recorded as the time between key-presses. After the final word of each item, a question appeared and participants responded by pressing a key. After an incorrect answer, the word "INCORRECT" flashed briefly on the screen. Approximately half of the comprehension questions were correctly answered as "yes" and half "no". Before the main experiment, a short list of practice items and questions was presented in order to familiarise the participant with the task.

Results

The data from two participants were not included in the following analyses. These participants answered only 52% and 64% of the comprehension questions from this experiment correctly. The next lowest accuracy rate was 69%.

Comprehension performance. Overall, participants understood the experimental sentences well, as evidenced by a mean accuracy rate of 87% on the comprehension questions. Figure 1 shows the percentage of comprehension questions that were answered correctly for each condition in the experiment.

A 3 $\times$ 3 repeated measures ANOVA on the question-answering data showed a main effect of NP2 type, $F_1(2, 78) = 3.4$, MSE = .022, $p < .05$; $F_2(2, 70) = 3.1$, MSE= .021, $p = .05$, which was likely due to the higher mean accuracy when NP2 was a pronoun than when it was not. An interaction between NP1 type and NP2 type was also reliable, $F_1(4, 156) = 6.9$, MSE = .023, $p < .001$; $F_2(4, 140) = 3.9$, MSE = .035, $p = .005$, likely

[2] The software was Linger 1.7 by Doug Rohde.

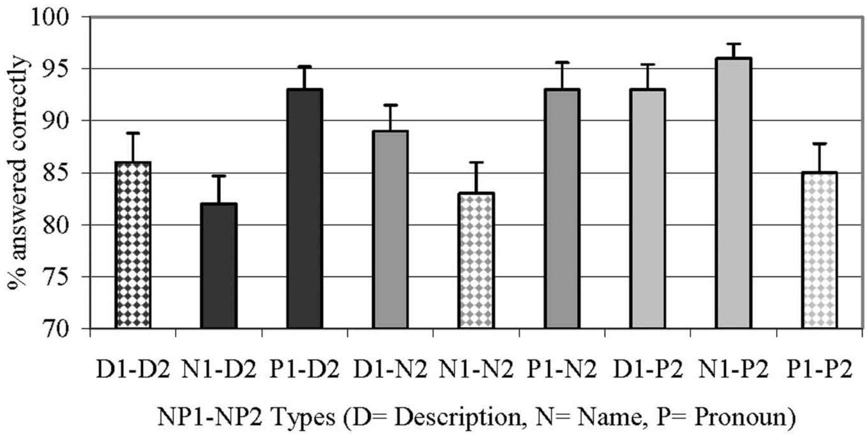

Figure 1. Percentage of comprehension questions answered correctly by condition.

due to a reliable NP match effect. Questions about sentences with non-matching NPs were answered reliably more accurately than sentences with matching NPs, $F_1(1, 39) = 12.8$, MSE $= .007$, $p = .001$; $F_2(1, 35) = 9.3$, MSE $= .008$, $p < .005$.

Reading times. Reading times beyond 3 standard deviations from the mean for a given word position in a given condition were replaced with the mean, affecting less than 2% of the data. Figure 2 reports average reading times per word across the sentence, averaged over conditions sharing the same type of NP2. Figure 3 reports reading times on the critical word. A 3×3 repeated measures ANOVA at the critical word—the verb—showed a reliable effect of NP2 type, as predicted by the referential hypothesis, because pronoun conditions averaged 30 ms faster than name conditions, which averaged 28 ms faster than description conditions, $F_1(2, 78) = 23.3$, MSE $= 4315$, $p < .001$; $F_2(2, 70) = 10.7$, MSE $= 8000$, $p < .001$. All pairwise comparisons among these means were reliable (ps $< .05$). This main effect of NP2 was also present on each of the two words following the verb, verb + 1: $F_1(2, 78) = 8.8$, MSE $= 1718$, $p < .001$; $F_2(2, 70) = 3.3$, MSE $= 4232$, $p < .05$; verb + 2: $F_1(2, 78) = 7.9$, MSE $= 1195$, $p < .001$; $F_2(2, 70) = 4.0$, MSE $= 2025$, $p < .05$.

At the verb, there was no effect of NP1 type, and an interaction between NP1 type and NP2 type was marginally reliable only in the analysis against participant variability, $F_1(4, 156) = 2.2$, MSE $= 3504$, $p = .08$; $F_2 < 2$, $p >$ .3. This weak effect reflected a match effect that was also only reliable by participants, such that the mean of the three conditions with matching NP types was 14 ms slower than the mean of the six conditions with different NP types, $F_1(1, 39) = 6.2$, MSE $= 630$, $p < .05$; $F_2 = 2.1$, $p > .1$. The

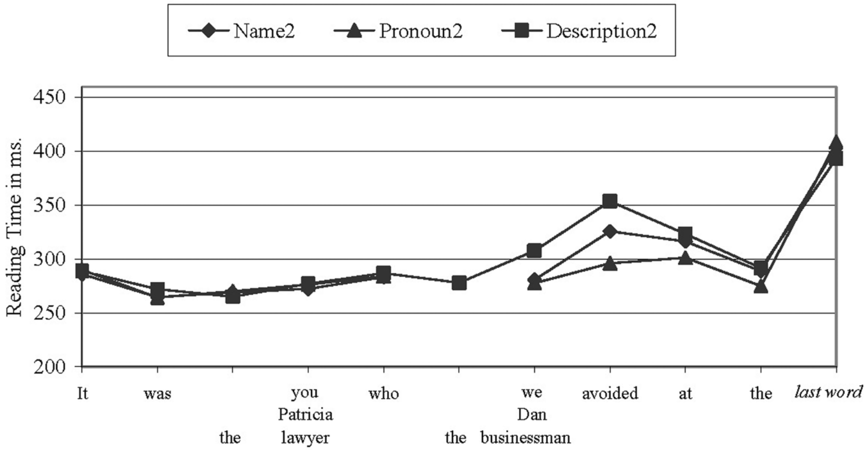

Figure 2. Average reading time per word averaged over conditions with the same NP type in NP2 position. Because the number of words in the adjunct phrase varied, we include reading times at the two words following the verb and then reading times on the sentence final word.

weakness of this effect is likely due to the fact that while the name-name and description-description conditions were read among the slowest on the verb, the pronoun-pronoun condition was among the fastest. This pattern continued until sentence end, when the pronoun-pronoun condition showed slow reading times similar to the other matching NP conditions. Figure 4 reports average reading times per word across the sentence, with the matched conditions highlighted. Figure 5 reports reading times on the sentence-final word. At the sentence-final word there was a reliable

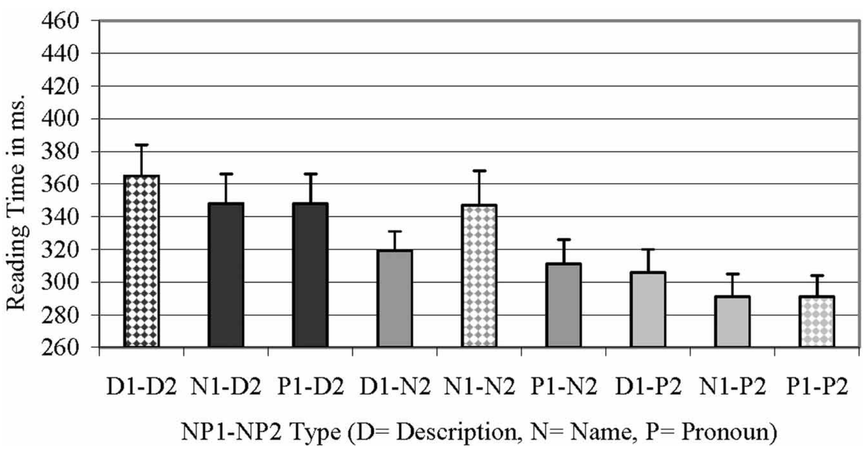

Figure 3. Average reading time on the verb in ms, by condition.

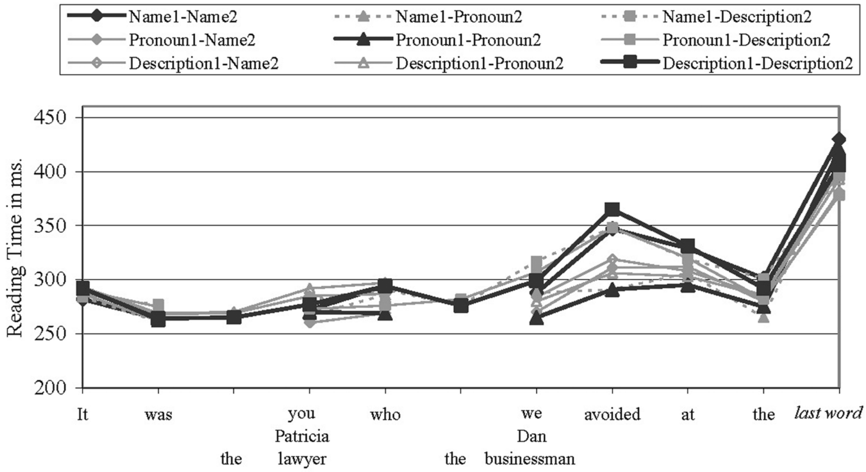

Figure 4. Average reading times per word with matched conditions highlighted. Because the number of words in the adjunct phrase varied, we include reading times at the two words following the verb and then reading times on the sentence final word.

interaction between NP1 and NP2, $F_1(4, 156) = 3.4$, MSE = 4178, $p = .01$; $F_2(4, 140) = 2.7$, MSE = 5507, $p < .05$, but no main effects. This interaction reflected a reliable NP match effect, as matching conditions were read 25 ms slower than non-matching conditions, $F_1(1, 39) = 10.0$, MSE = 1287, $p < .005$; $F_2(1, 35) = 7.8$, MSE = 1565, $p < .01$.

The fact that the critical word immediately followed NP2 in these stimuli is a potential concern, because the pattern of reading times on the critical

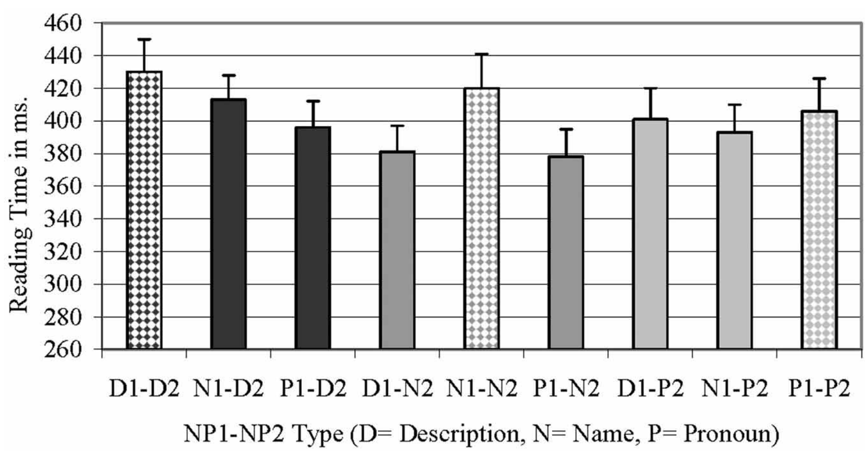

Figure 5. Average reading time on the last word of the sentence in ms, by condition

word may have been affected by spill-over of lexical access effects or referential access effects that began on NP2. To evaluate this possibility, we investigated reading times on the word after NP1 because the pronouns, names and descriptions used in both NP positions were similar. If spill-over from lexical access or referential access was a problem, it should have affected times after NP1 and NP2 similarly. Individual comparisons showed that the average reading time on the word following NP1 was faster when NP1 was a pronoun than when it was a name or description (ps $< .01$), following the same pattern as after NP2. However, there was only a two ms reading time difference on the word after NP1 when NP1 was a name vs. description (ps $> .6$). Therefore, while reading times at the critical word in the pronoun condition may have been affected by spillover from the processing of the previous word, the difference between the name and description NP2 conditions at the verb is not likely due to spill-over from processing that began on the NP.

DISCUSSION

The results confirm predictions of both the referential accessibility and similarity based hypotheses, but suggest that their effects may have different time courses. At the verb, the referential accessibility hypothesis' prediction that description conditions would be read slowest, name conditions faster and pronoun conditions fastest was fully supported, but there was only mixed evidence for similarity based interference, as the description-description and name-name conditions were read slowly but the pronoun-pronoun condition were read extremely quickly. However, comprehension question accuracy and sentence final reading times showed strong effects of similarity based interference, with low accuracy and slow times for conditions with matching NPs, but no effects of NP2 type besides a slight increase in accuracy when NP2 was a pronoun.

The on-line results support Warren and Gibson's (2002) referential accessibility hypothesis that during structure building, it is easier to integrate across pronouns than names and across names than descriptions. Later measures associated with wrap-up processes and proposition retrieval provide the strongest evidence for Gordon et al.'s (2001) similarity based integration cost hypothesis. One possible explanation for this pattern is that the referential-to-syntactic interference described in Warren and Gibson (2002) occurs during structure building and has no residual effects once the integrations are constructed. Gordon et al.'s similarity based interference, on the other hand, may primarily affect the process of consolidating a conceptual representation and retrieving it from memory. This suggestion is not incompatible with Gordon et al.'s (2001) finding of interference effects at the verb, because average reading times in that study were approximately twice as long as in this study, perhaps because at the verb their participants may have initiated the kinds of processing that our participants performed off-line. The different time courses of the effects in the current study suggest that working memory load due to accessing or accommodating referents may affect what Caplan and Waters (1999) call interpretive processing while similarity based interference may affect either both interpretive and post-interpretive processing or possibly only post-interpretive processing.

One possible reason that the pronoun-pronoun condition did not show similarity based interference at the critical word may be that first person pronouns are explicitly marked for case. Case-marking provides memory cues as to which nouns are subjects and objects, making interference less of a problem when assigning roles to NPs. If the reduced processing difficulty in the pronoun-pronoun condition was due to case-marking, then the results are entirely consistent with similarity based interference playing a role in the on-line building of sentence representations. Note, however, that the case marking of indexical pronouns cannot be the only reason integrations involving indexical pronouns are easier, because Gordon et al. (2001) found an advantage for pronouns when using only the non-case-marked ‘‘you’’. Interestingly, the fact that the pronoun-pronoun conditions patterned with the name-name and description-description conditions in the wrap-up and retrieval measures but not on the critical word suggests that the processing difficulty apparent during consolidation and retrieval is not just a consequence of previous processing difficulty. Rather, this pattern suggests that NPs with the same type interfere at the level of a conceptual representation, independent of whether they interfered during on-line structure building.

Whereas the current experiment tested factors thought to affect retrieval difficulty in sentence comprehension, other factors have been shown to influence processing difficulty in RCs. For example, recent results suggest that processing difficulty associated with object-extracted RCs is moderated or eliminated when the head noun is inanimate (Mak, Vonk, & Schriefers, 2002; Traxler, Morris, & Seely, 2002). These findings do not follow from a referential accessibility based theory. Animacy could naturally fit into a similarity theory, as a factor affecting the similarity of two NPs. However, the animacy results do not support a similarity based theory based on simple confusability in memory, as noted by Traxler et al. (2002). Such a theory would predict that the difficulty of retrieving an NP from a set in memory including one animate and one inanimate NP should not be affected by the order in which those NPs appeared in the sentence. Traxler et al. (2002) showed an asymmetry in difficulty between RCs with one animate and one inanimate NP dependent on their positions within the RC, inconsistent with a simple confusability theory. These findings suggest that the animacy of the head noun is an independent factor influencing RC processing difficulty. One possible reason animacy might affect processing difficulty is that inanimate subjects are likely generally less plausible than animate subjects. People may be sensitive to such a difference either as a result of experience with events in the world and/or the language they are exposed to describing those events (MacDonald, 1999).

The results of the current experiment suggest that multiple factors influence memory load during the processing of linguistic dependencies. One of these factors is the similarity of the NPs that initiate dependencies. This supports the claim that similarity based interference affects linguistic processing. Another factor is the givenness status of the NPs between the endpoints of a dependency. This suggests that increasing the referential processing required between the endpoints of a dependency impedes the retrieval of the first endpoint. While Gibson (1998, 2000) suggested that this increased difficulty is due to decay, in light of the evidence for similarity based interference it may be more parsimonious to assume it is due to retrieval interference. The current experiment suggests that theories of linguistic complexity must be modified in order to take both similarity based interference and givenness into account when quantifying dependency length.

Manuscript received July 2004
Revised manuscript received January 2005

REFERENCES

Ariel, M. (1990). *Accessing noun-phrase antecedents*. London: Routledge.

Bever, T. G. (1970). The cognitive basis of linguistic structures. In J. R. Hayes (Ed.), *Cognition and the development of language* (pp. 279–362). New York: John Wiley & Sons.

Caplan, D., & Waters, G. S. (1999). Verbal working memory and sentence comprehension. *Brain and Behavioral Sciences*, *22*, 77–126.

Chomsky, N., & Miller, G. A. (1963). Introduction to the formal analysis of natural languages. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (Vol. 2, pp. 269–321). New York: John Wiley & Sons.

Desmet, T., Brysbaert, M., & De Baecke, C. (2002). The correspondence between sentence production and corpus frequencies in modifier attachment. *Quarterly Journal of Experimental Psychology: Section A*, *55(3)*, 879–896.

Ford, M. (1983). A method for obtaining measures of local parsing complexity throughout sentences. *Journal of Verbal Learning and Verbal Behavior*, *22*, 203–218.

Frazier, L. (1987). Sentence processing: A tutorial review. In M. Coltheart (Ed.), *Attention and Performance XII*. Hove, UK: Lawrence Erlbaum Associates Ltd.

Gibson, E. (1991). *A computational theory of human linguistic processing: memory limitations and processing breakdown*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA.

Gibson, E. (1998). Linguistic complexity: Locality of syntactic dependencies. *Cognition*, *68*, 1–76.

Gibson, E. (2000). The dependency locality theory: A distance-based theory of linguistic complexity. In Y. Miyashita, A. Marantz, & W. O'Neil, (Eds.), *Image, language, brain* (pp. 95–126). Cambridge, MA: MIT Press.

Gibson, E., & Pearlmutter, N. (1998). Constraints on sentence comprehension. *Trends in Cognitive Science*, *2*, 262–268.

Gordon, P. C., Hendrick, R., & Johnson, M. (2004). Effects of noun phrase type on sentence complexity. *Journal of Memory and Language*, *51*, 97–114.

Gordon, P. C., Hendrick, R., & Johnson, M. (2001). Memory interference during language processing. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *27*, 1411–1423.

Grodner, D., & Gibson, E. (2005). Consequences of the serial nature of linguistic input for sentential complexity. *Cognitive Science*, *29*, 261–290.

Gundel, J., Hedberg, H., & Zacharski, R. (1993). Referring expressions in discourse. *Language*, *69*, 274–307.

Hakes, B., Evans, J., & Brannon, L. (1976). Understanding sentences with relative clauses. *Memory and Cognition*, 4, 283–296.

Haviland, S. E., & Clark, H. H. (1974). What's new? Acquiring new information as a process of comprehension. *Journal of Verbal Learning and Verbal Behavior*, *13*, 512–521.

Hawkins, J. A. (1994). *A performance theory of order and constituency*. Cambridge: Cambridge University Press.

Holmes, V. M. & O''Regan, J. K. (1981). Eye fixation patterns during the reading of relative clause sentences. *Journal of Verbal Learning and Verbal Behavior*, *20*, 417–430.

Just, M. A., Carpenter, P. A., & Woolley, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General*, *111*, 228–238.

Kimball, J. (1973). Seven principles of surface structure parsing in natural language. *Cognition*, *2*, 15–47.

King, J., & Just, M. A. (1991). Individual differences in syntactic processing: The role of working memory. *Journal of Memory and Language*, *30*, 580–602.

Mak, W. M., Vonk, W., & Schriefers, H. (2002). The influence of animacy on relative clause processing. *Journal of Memory and Language*, *47*, 50–68.

Lewis, R. L. (1996). Interference in short term memory: The magical number two (or three) in sentence processing. *Journal of Psycholinguistic Research*, *25*, 93–117.

MacDonald, M. (1999). Distributional information in language comprehension, production and acquisition: Three puzzles and a moral. In MacWhinney, Brian (Ed.), *The emergence of language.* Mahwah, NJ: Lawrence Erlbaum Inc.

MacDonald, M., & Christiansen, M. (2002). Reassessing working memory: Comment on Just and Carpenter (1992) and Waters and Caplan (1996). *Psychological Review*, *109* (*1*), 35–54.

MacWhinney, B., & Pleh, C. (1988). The processing of restrictive relative clauses in Hungarian. *Cognition*, *29*, 95–141.

Murphy, G. L. (1984). Establishing and accessing referents in discourse. *Memory and Cognition*, *12*(*5*), 489–497.

Pearlmutter, N., & Gibson, E. (2001). Recency in verb phrase attachment. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *27*, 574–590.

Tanenhaus, M. K., & Trueswell, J. C. (1995). Sentence comprehension. In J. Miller & P. Eimas (Eds.), *Speech, language, and communication.* London: Academic Press.

Traxler, M., Morris, R., & Seely, R. (2002). Processing subject and object extracted relative clauses: Evidence from eye-movements. *Journal of Memory and Language*, *47*, 69–90.

Van Dyke, J. A., & Lewis, R. L., (2003). Distinguishing effects of structure and decay in attachment and repair: A cue-based parsing account of recovery from misanalyzed ambiguities. *Journal of Memory and Language*, *49*, 285–316.

Wanner, E., & Maratsos, M. (1978). An ATN approach in comprehension. In M. Halle, J. Bresnan, & G. Miller, (Eds.), *Linguistic theory and psychological reality* (pp. 119–161). Cambridge, MA: MIT Press.

Warren, T., & Gibson, E. (2002). The influence of referential processing on sentence complexity. *Cognition*, *85*, 79–112.

Waters, G. S., Caplan, D., & Hildebrandt, N. (1987). Working memory and written sentence comprehension. In Coltheart, M. (Ed.), *Attention and performance: Vol. XII: The psychology of reading* (pp. 531–555). Hove, UK: Lawrence Erlbaum Associates Ltd.

Webber, B. L. (1979). *A formal approach to discourse anaphora*. New York: Garland Publishing.

APPENDIX

Experimental Items

1. It was {the physicist/John/us} who {the dean/Bruce/I} praised at the meeting.
2. It was {the lawyer/Patricia/you} who {the businessman/Dan/we} avoided at the party.
3. It was {the politician/Ann/me} who {the student/Susan/you} criticised at the dinner.
4. It was {the fisherman/Jody/you} who {the clerk/Elizabeth/I} saw at the store.
5. It was {the banker/Sue/us} who {the lobbyist/Dee/you} praised in front of the crowd.
6. It was {the dancer/Jill/you} who {the reporter/Rose/I} phoned on New Year's Eve.
7. It was {the architect/Ted/you} who {the fireman/Wes/we} liked before the argument began.
8. It was {the detective/Jack/you} who {the secretary/Bill/I} disliked during card games.
9. It was {the politician/Luke/us} who {the mailman/Mark/you} insulted after reading the newspaper article.
10. It was {the governor/Barb/us} who {the comedian/Gwen/I} complimented in the fancy restaurant.
11. It was {the actor/Kim/me} who {the director/Fay/we} thanked before the show.
12. It was {the busboy/Pete/you} who {the cashier/Nick/I} distrusted after the restaurant closed.
13. It was {the violinist/Bob/us} who {the conductor/Max/you} complimented at Carnegie-Hall.
14. It was {the tutor/Todd/us} who {the student/Brad/I} questioned during summer vacation.
15. It was {the editor/Pam/me} who {the author/Jen/we} recommended after a new merger was announced.
16. It was {the tailor/Reid/me} who {the customer/Kate/you} described at the banquet.
17. It was {the admiral/Ken/us} who {the general/Jim/I} advised before the trip got underway.
18. It was {the referee/Joy/me} who {the coach/Eve/we} blamed for the outcome of the game.
19. It was {the lawyer/Seth/me} who {the auditor/Greg/you} interviewed in the very small office.
20. It was {the plumber/Lynn/you} who {the electrician/Beth/I} called from the pay-phone.
21. It was {the clown/Liz/us} who {the magician/Meg/you} entertained in the auditorium.
22. It was {the child/Nate/us} who {the traveler/Doug/I} protected in the burning building.
23. It was {the gardener/Dawn/you} who {the homeowner/Fran/we} envied after the lottery ended.
24. It was {the director/Paul/us} who {the actress/Lisa/you} impressed at the audition.
25. It was {the journalist/Nancy/you} who {the mayor/Bob/I} greeted at the door.
26. It was {the tenant/Abby/you} who {the landlord/Steve/we} sued for a lot of money.
27. It was {the housewife/Laura/me} who {the salesman/Max/you} startled at the forum.

28. It was {the plumber/Jen/us} who {the electrician/Mary/I} complimented on the excellent repair job.
29. It was {the nurse/Lee/you} who {the scientist/Pam/we} recognized at the coffee shop.
30. It was {the secretary/Jon/us} who {the editor/Kathy/you} hired some time last week.
31. It was {the consultant/Mary/me} who {the administrator/Christina/we} selected as the spokesperson for the company.
32. It was {the technician/Martin/me} who {the executive/Allen/we} acknowledged for good work at the awards ceremony.
33. It was {the witness/Ellen/me} who {the bodyguard/Max/we} protected before the important trial.
34. It was {the engineer/Katie/me} who {the customer/Joe/we} recognised as the most knowledgeable about the product.
35. It was {the teacher/Lauren/me} who {the secretary/Taylor/you} offended with the vulgar joke at the movie theater.
36. It was {the coach/Mike/me} who {the lawyer/Morgan/you} saw in the store with the party supplies.