Memory Retrieval Effects on Filler-Gap Processing

Philip Hofmeister (philiph@stanford.edu)

Department of Linguistics
Stanford University
Stanford, CA 94305

Abstract

In this paper, I argue that the retrieval of filler-phrases in syntactic filler-gap dependencies is facilitated by encoding more information in the filler-phrase. Three self-paced reading experiments provide the evidence for this memory facilitation hypothesis: reading times decrease when more explicit *wh*-phrases and indefinites occur as fillers in three different syntactic constructions. Crucially, the decrease in reading times becomes visible only at the point of memory retrieval, i.e. the gap site. There is no effect of the amount of information encoded in the filler-phrase prior to the retrieval site.

Keywords: filler-gap dependencies; memory retrieval; sentence processing; syntactic islands

Introduction

The processing of so-called filler-gap dependencies, as in (1), is known to be affected by a variety of factors. In these constructions, a linguistic constituent (the filler) is displaced from its typical position, leaving an empty syntactic position (the gap). As Hawkins (1999) points out, there are difficulties in identifying the correct gap site, since it has no overt manifestation, and while trying to identify the gap site, the filler must be held in working memory while other linguistic constituents are processed along the path from filler to gap. The difficulty of processing these constructions further increases with distance (Gibson, 1998), interference effects (Gordon, Hendrick, & Johnson, 2004; Van Dyke & McElree, 2006), as well as processing load imposed by other referential entities along the filler-gap path (Warren & Gibson, 2002).

(1) [Which cup] did you drink from ___ before you started feeling ill?

Here, I investigate how the amount of information encoded in a filler-phrase affects filler-gap dependency processing. On the extreme view that storing more material always incurs more processing costs, a more complex and informative filler should impair dependency processing. I propose, however, a type of memory facilitation hypothesis that says that linguistic elements that encode *more* information (including syntactic, semantic, and lexical information) facilitate their own retrieval later. In other words, richer linguistic descriptions of discourse entities make recall of these entities easier (all else being equal).

Previous work in psychology, psycholinguistics, and pragmatics is suggestive of the effect that informativity can have on memory and discourse. For instance, Bradshaw and Anderson (1982) provide evidence that sentence recall improves when presentation of a sentence occurs along with other causally-related propositions. To explain these results, Bradshaw and Anderson suggest that elaborating on a given topic increases the number of possible retrieval paths to a given memory trace. Having more retrieval paths, in turn, implies

higher chances of successful recall. Also, various theoretical works in linguistics (Pesetsky, 1987, 2000; Deane, 1991; Kluender, 1992; Frazier & Clifton, 2002) have noticed that the nature of the filler-phrase impacts acceptability in a significant way for certain extraction types. For instance, Pesetsky (1987) observes that, in multiple *wh*-questions when the object *wh*-phrase is extracted over the subject *wh*-phrase (also known as "Superiority-violations" (Chomsky, 1973)), acceptability improves when informationally richer *which*-phrases are used (2a), as opposed to bare *wh*-items like *who* and *what* (2b):

(2) a. Which treatment did which patient receive? >

     b. What did who receive?

Pesetsky explains this contrast as the result of a distinct grammatical mechanism that he terms 'D(iscourse)-linking'.[1] Similarly, Kluender (1992) discusses how increased "referential specificity" facilitates "long *wh*-movement." For the most part, however, these accounts are purely theoretical without any supporting empirical evidence obtained through controlled experimentation and concentrate exclusively on acceptability (cf. Frazier and Clifton (2002)). More importantly, this literature offers few explanations for why informationally richer fillers should improve filler-gap dependencies.

In this paper, I present the results of three self-paced reading studies that all provide positive evidence for the hypothesis that more informative linguistic descriptions lead to faster recall. In each experiment, the results indicate that reading times are significantly reduced at the key point of memory retrieval in a filler-gap dependency (as well as in the spillover region)[2] when a more explicit filler phrase is used.

The first study considers the effect of *wh*-phrase informativity on the processing of so-called syntactic islands. These constructions involve the extraction or dislocation of a linguistic element across a syntactic boundary which supposedly blocks the extraction of the element (Ross, 1967). One particular constraint, the *wh*-island constraint, says that a *wh*-phrase cannot be extracted out of an embedded clause with another

---

[1] I avoid the use of the term "D-linked" in this paper to characterize certain *wh*-phrases. No precise formulation of or motivation for the D-linking analysis has ever been offered, and there is no reliable method for ascertaining whether a given phrase should be D-linked or not. Indeed, Pesetsky himself notes that the division between scope-taking and D-linked *wh*-expressions is not absolute, admitting that D-linking (as measured by the ability to participate in superiority-violations) can occur even with the less informative bare *wh*-items like *who* and *what*.

[2] Spill-over results from continued processing of a particular lexical item or region into subsequent regions. Since most language processing tasks typically are followed by another, the completion of one task (especially complex and challenging processing tasks) often overlaps with the processing of one or more subsequent tasks.

wh-phrase in complementizer position. Such a constraint acts as an explanation for the markedness of the questions below:

(3) Sawyer said who committed this atrocity.
     What did Sawyer say [$_S$ who committed ___] ?

(4) The journalist wondered whether the president planned the attack.
     What did the journalist wonder [$_S$ whether the president planned ___] ?

While traditional syntactic accounts label them as ungrammatical, a number of sources trace the acceptability of structures involving islands to the processing complexity of these constructions (Alexopoulou & Keller, 2003; Kluender, 1998; Hofmeister, Jaeger, Sag, Arnon, & Snider, 2007). Given the high processing costs associated with these constructions, therefore, they constitute an optimal setting for testing the predictions of the memory facilitation hypothesis. That is, simple, unary *wh*-questions are presumably processed without much difficulty and thus memory retrieval processes may be operating at optimal efficiency. In a given environment, if retrieval can take place in the minimum amount of time that the cognitive architecture will allow, then any strategy meant to facilitate retrieval will not be able to lower retrieval times any further. In contrast, *wh*-islands present a suitably challenging setting for language processing, making it more likely that advantages for sentence processing will be visible.

To allay concerns that the observed effect in *wh*-island constructions ultimately is artefactual, a second study is reported here that involves nested dependencies. These constructions, involving one filler-gap dependency nested inside a larger dependency, also impose significant processing difficulties on the comprehender; however, no one seriously questions the grammaticality of these constructions. Lastly, I describe here the results of an experiment that evaluates how informativity affects the processing of clefted indefinites, e.g. *It was a neurosurgeon he consulted on Friday*. This study complements the other studies by showing that the observed effects are not limited to *wh*-dependencies. I conclude this paper with some observations on how this data accords with other findings about the choice of NP form and the potential underlying causes of this facilitation.

Experiment 1: *Wh*-Islands

Methodology   In self-paced reading time experiments, subjects read sentences at their own pace. Initially, they are presented with a screen of dashes separated by spaces, representing the words for that experimental item. With each press of a predefined key, a new word appears on the screen and the previous word disappears. Following the initial sentence, the subjects' comprehension is tested with a follow-up question. Subjects are informed beforehand that they are participating in a reading comprehension survey and that their ability to understand small passages of English is being tested. Reading times for each word are recorded, collected, and analyzed as a measure of processing difficulty, with higher reading times representing more processing difficulty. Reaction times and question-answer accuracy were also recorded.

Materials   All experimental items consisted of *wh*-islands, presented as main clause interrogatives, as in (5). Subjects initially saw a declarative sentence, after which a comprehension question was presented. In this experiment and the next, the real stimuli were the comprehension questions themselves; the initial context sentences merely justified the presence of these questions which would be entirely unnatural without any background. The *wh*-phrase always corresponded to the animate object of the embedded verb. The stimuli varied in terms of whether the sentence-initial *wh*-phrase was the bare *wh*-item *who* (BARE) or the more complex and informative *which*-phrase (WHICH). Additionally, a third condition was included to serve as a baseline against which the results could be compared. This condition differed only in terms of the complementizer and because this complementizer was always *that*, these questions do not represent instances of syntactic islands.

Twenty-four experimental items and forty-eight fillers were included in this study. Twelve of the fillers were *wh*-islands, such that half of the overall items were *wh*-islands and the other half were not syntactic islands. Furthermore, an equal number of questions began with *who*, *what*, and *which*-phrases.

(5) Albert learned that the managers dismissed the employee with poor sales after the annual performance review.

     BARE: Who did Albert learn whether they dismissed after the annual performance review?
     WHICH: Which employee did Albert learn whether they dismissed after the annual performance review?
     BASELINE: Who did Albert learn that they dismissed after the annual performance review?

After reading the question, subjects answered the question by selecting from multiple choices. Of the three possible answers presented to them, one was correct, another was lexically and syntactically similar, and the third option differed drastically. Only correctly answered stimuli are considered here.

Eight practice examples preceded the real experimental items. Each participant saw each item in exactly one condition (Latin-square design). Residual reading times were derived for each word on the basis of a linear regression equation that computes reading time for each individual as a function of word length. This method effectively reduces variability due to individual differences in reading times. Where I report results, only the residual reading times are considered, rather than the raw reading times. The results from all experiments covered here were analyzed with repeated measures ANOVAs.

Participants   Twenty subjects participated in this reading-time experiment. All participants were Stanford University undergraduates who received course credit for their participation. The experiment was run immediately before or after another reading time experiment on multiple *wh*-questions. On average, each experimental session lasted about 45 minutes.
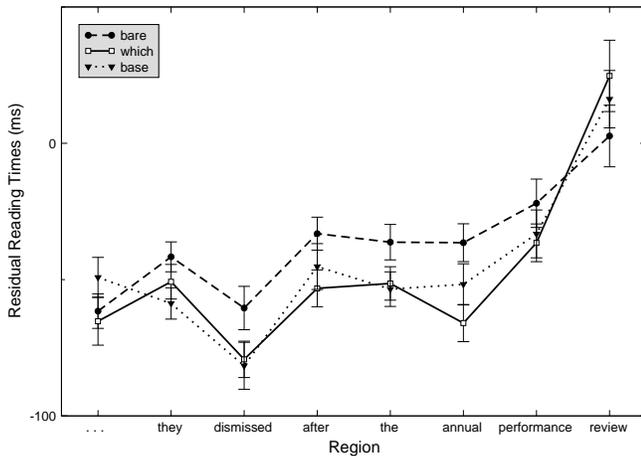
Figure 1: Residual reading times in experiment 1. The first region corresponds to the averaged reading times for the four words between the end of the *wh*-phrase, and up to the embedded subject, e.g. *did Albert learn whether*.

Results The results verify that more explicit *wh*-phrases lead to faster reading times at the subcategorizing verb ($F_{1}(1,19)$ = 3.909, p = .063, $F_{2}(1,23)$ = 5.583, p <.05). When spillover regions are considered, the effect of phrase informativity is even more observable and statistically significant. Considering the verb and the subsequent modifying material together, the difference between *which*-N' phrases and *who* is significant by both subjects and items ($F_{1}(1,19)$ = 7.919, p <.01, $F_{2}(1, 23)$ = 10.504, p <.01). Reading times on the word immediately following the verb (where the presence of the gap is confirmed) alone reflect significantly faster reading times for the WHICH condition ($F_{1}(1,19)$ = 7.00, p = .016, $F_{2}(1,23)$ = 5.412, p = .029) . Indeed, this trend continues for three words after the verb, as pictured in Figure 1, such that reading times in the bare-*wh* condition remain slower until sentence-final wrap-up effects obscure the difference between conditions. This suggests that at least some (if not a large proportion) of the difficulty associated with retrieving the *wh*-object continues after the verb is processed. The processing facilitation associated with the more informative *which*-N' phrases also eliminated any substantial difference between the island-violating *which*-condition and the baseline condition, which does not violate any supposed constraints on extraction.

Crucially, the effect of informativity first emerges at the retrieval site, i.e. the subcategorizing verb. In the regions preceding the verb that includes the matrix subject and verb, as well as the complementizer and embedded subject, there is no significant difference across conditions. In other words, the processing advantage due to increased informativity of a stored constituent appears at the retrieval and integration site and not during the storage interim. Hence, this evidence argues in favor of the hypothesis that encoding more information about a discourse entity can facilitate memory retrieval of that entity.

Reaction times to the question showed a strong advantage

for the WHICH condition over the BARE condition ($F_{1}(1,19)$ = 13.664, p = .002, $F_{2}(1,23)$ = 10.778, p = .003). The main effect of *wh*-phrase type on reaction times, however, is confounded by the fact that two of the multiple choice answers contained the noun appearing in the *which*-phrase. This leads to the possibility that the reaction time results underlyingly stem from a priming or word recognition effect. Question-answer accuracy did not differ significantly across conditions.

Experiment 2: Nested Dependencies

In order to counter any objections that the above results merely reflect an atypical processing pattern associated with the comprehension of island constructions, it is necessary to consider the same effect in another complex, but unquestionably grammatical context. This experiment expands upon the previous study by considering another challenging filler-gap dependency, but one which has not been previously labeled ungrammatical.

Materials The twenty experimental items in this study were all nested dependency interrogatives of the sort illustrated in (6). The interrogatives varied in terms of whether the fronted, inanimate *wh*-phrase was a bare *wh*-item (SIMPLE), a *which*-N' phrase (WHICH), or a *which*-N' phrase with an additional adjective (COMPLEX). These interrogatives acted as the comprehension questions for preceding sentences, as in the previously described study.

(6)   Scooter hid from the reporter who talked about the recent ABC political poll on a recent evening news segment.

SIMPLE: What did the reporter that Scooter avoided discuss during an evening news segment?
WHICH: Which poll did the reporter that Scooter avoided discuss during an evening news segment?
COMPLEX: Which political poll did the reporter that Scooter avoided discuss during an evening news segment?

Eighty-eight fillers were included in the item list, including an equal number of subject relatives. Across each experimental list, subjects saw an equal number of questions beginning with *who*, *what*, and *which*.

Participants Thirty-five native English-speaking Stanford University undergraduates participated in this study to fulfill a course requirement in a linguistics class. The subjects had not participated in the previous experiment on *wh*-islands.

Results As predicted, the more informative *wh*-expressions led to faster reading times at the subcategorizing verb. The "complex" *which*-phrase containing an adjective produced faster reading times at the verb than the bare *wh*-phrase ($F_{1}(1,34)$ = 5.158, p <.05; $F_{2}(1,19)$ = 5.260, p <.05). Similarly, the use of the "simple" *which*-phrase without the adjective also resulted in faster reading times at the critical region of the verb ($F_{1}(1,34)$ = 7.684, p <.01; $F_{2}(1,19)$ =
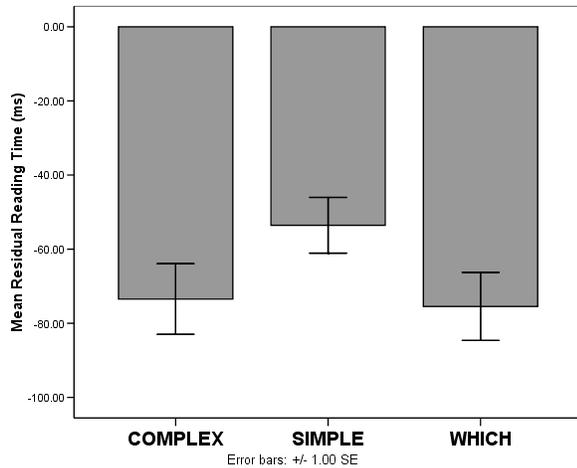
Figure 2: Residual reading times at matrix verb (e.g. *discussed*) in experiment 2

Figure 3: Residual reading times at first word after matrix verb (spillover) in experiment 2

7.81, p = .012). This effect intensifies during the immediately following spillover region (COMPLEX vs. SIMPLE: $F_1(1,34) = 15.520$, p <.001, $F_2(1,19) = 4.914$, p <.05; WHICH vs. SIMPLE: $F_1(1,34) = 5.562$, p <.05, $F_2(1,19) = 5.229$, p <.05). As is evident from Figure 2, however, there is no significant difference between the WHICH and COMPLEX conditions. Presumably, the addition of one extra adjective was insufficient to create a noticeable difference in retrievability. Reaction times reflect a speed-up for the COMPLEX condition as compared to the SIMPLE condition ($F_1(1,34) = 6.44$, p = .016, $F_2(1,19) = 12.26$, p <.01). There was also a non-significant trend for faster reaction times in the COMPLEX condition as compared to the WHICH condition ($F_1(1,34) = 3.42$, p<.1, $F_2(1,19) = 2.35$, p>.1). As in the previous experiment, the reaction time results are confounded by the fact that the *which*-N' phrases contained the same head noun as the correct answer. There were no significant differences in question-answer accuracy across conditions.

The factor of informativity once again does not significantly impact reading times until retrieval of the stored filler-phrase is necessary, in line with the predictions of a retrieval-based account. Reading times prior to the two verbs show no effect of the amount of information encoded in the *wh*-phrases.[3] Interestingly, however, at the verb (*avoided* in (6) above) immediately preceding the actual subcategorizing verb, this effect is evident: the WHICH and COMPLEX conditions result in faster reading times than SIMPLE. This result, though, is unsurprising given a memory retrieval account, since processing of this verb also requires retrieval and a *wh*-phrase is still being held in memory, potentially interfering with retrieval of the subject NP.

---

[3] This absence of an effect prior to the verb or retrieval sites also argues against any alternative explanation of the facts that subjects read the *which*-conditions faster strictly due to an early anticipation of the correct answer.
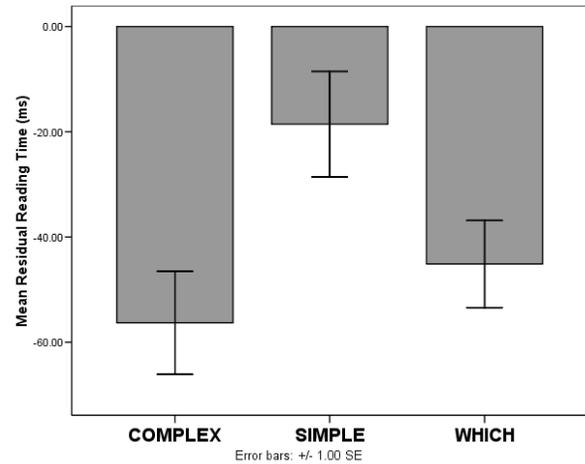
Experiment 3: Clefted Indefinites

The preceding two experiments concentrate exclusively on filler-gap dependencies involving *wh*-phrases. The proposal made here, however, should cover other types of filler-gap dependencies as well. In other words, the more information encoded in any type of filler-phrase, the faster memory retrieval should happen (again, all else being equal). To test this prediction, I consider filler-gap dependencies involving indefinite NPs of varying levels of informativity.

Materials Sixteen clefted indefinites constituted the experimental data set, as depicted in (7). Conditions varied in terms of how many adjectives preceded the head noun: zero (SIMPLE), one (MID), or two (COMPLEX). In all items, the clefted indefinite was followed by a relative clause containing a five-word subject NP and then a transitive verb with an object gap, requiring the retrieval of the clefted indefinite phrase. In contrast to the other two experiments, these stimuli were followed by comprehension questions, rather than acting as the comprehension questions themselves. The comprehension questions for these sixteen items always asked about the identity of the individual referred to with the clefted indefinite. Other experimental fillers, however, involved clefts and comprehension questions asking about various other aspects of the sentence meaning besides the clefted phrase. Three multiple choice answers were provided after displaying the subsequent comprehension question.

(7) SIMPLE: It was a communist who the members of the club banned from ever entering the premises.
MID: It was an alleged communist who the members of the club banned from ever entering the premises.
COMPLEX: It was an alleged Venezuelan communist who the members of the club banned from ever entering the premises.

These experimental items were included as fillers in the aforementioned nested dependency experiment. Hence,
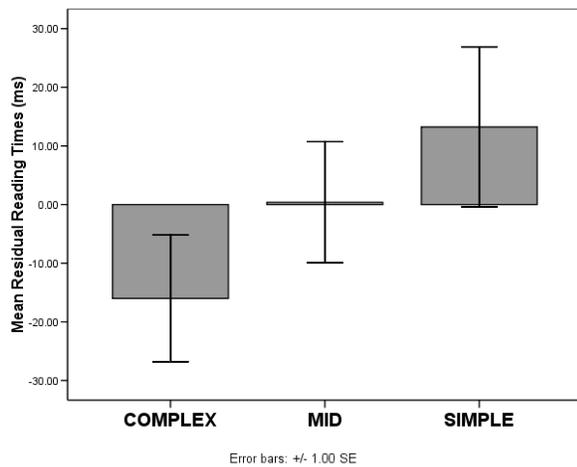
Figure 4: Residual reading times at subcategorizing verb in experiment 3

ninety-two other items acted as fillers for this data set.

Participants The same thirty-five native English speakers from experiment 2 participated in this experiment.

Results Following the results from the previous experiments, the most informative indefinites induced faster reading times at the verb than the least informative indefinites. This effect was nearly significant by subjects and significant by items ($F1(1,34) = 5.351$, $p = .063$, $F2(1,15) = 3.690$, $p < .05$). On the other hand, the comparisons between the SIMPLE and MID conditions and between the MID and COMPLEX conditions were not significant. The difference between these conditions is the presence or absence of one adjective, as was the case with the *which*-phrases in the preceding experiment. Hence, not just any single addition of information necessarily improves retrieval. Note, however, that this does not contradict the predictions of the memory facilitation hypothesis. This hypothesis does not claim or speculate about how much additional information is necessary in order to significantly improve memory retrieval.

As in the other experiments, the different levels of informativity in the filler phrase did not result in significantly different reading times prior to the verb where retrieval of the filler occurs. Spillover effects were not observed in this study, undoubtedly due to the relative processing ease of the sentences. Overall, the duplication of the facilitation in dependencies involving indefinites provides evidence, albeit limited, that the advantage of using more explicit filler-phrases is not restricted to *wh*-dependencies.

Discussion

Cumulatively, the three reading-time studies reflect reduced processing costs when more explicit fillers are used. It was found that *which*-phrases consistently led to faster processing times than bare-*wh*-phrases. Furthermore, indefinites encoded along with multiple adjectives facilitated processing, as compared to indefinites with no accompanying adjectives. Other experimental results extend the findings discussed here to other kinds of syntactic constructions, such as multiple *wh*-questions and so-called complex noun phrase violations (Hofmeister et al., 2007).

The lack of a significant difference between the two *which*-phrases in experiment 2 and between the MID indefinite condition and the other conditions likely stems from an insufficient difference in informativity. Notice that, while the difference between a *which*-N' phrase and *who* is also only one word, the former specifies much that the latter does not: in addition to specifying number information, the *which*-phrase further identifies the type of relevant individual or object. This goes beyond the specification of a mere attribute, as in the difference between *communist* and *alleged communist*.

Theoretically, additional information should also inhibit similarity effects of the sort that Gordon et al. (2004) observe. If interference effects are determined on the basis of various dimensions of similarity, e.g. definiteness, syntactic category, thematic role assignment, etc., additional lexical-semantic features should reduce the impact of interference from other candidates in memory. Indeed, the underlying cause for this facilitation may reside in the fact that adding or linking information to a mental representation causes it to be more distinctive from other potential competitors in memory. There are, of course, other possible explanations for this facilitation: the additional study time provided by extra morphophonemic material may allow more accurate and redundant encoding. Also, the addition of information to the representation of a discourse entity may increase the number of available retrieval paths in the sense of Bradshaw and Anderson (1982). In fact, these factors are not mutually exclusively and may additively contribute to the overall memory retrieval facilitation. In this sense, increased informativity may reflect an aggregate of dynamic factors. Future experimental work consequently seeks to identify the source of this sentence processing facilitation.

Several points suggest why the observed data should be regarded as a reflection of memory retrieval differences, instead of some other processing-based factor. As pointed out for all three experimental studies here, the effect of informativity appears only at the retrieval site. While integration also occurs at this point in sentence processing, there is little motivation for assuming that the integration of longer and more informative constituents is expedited. Furthermore, any explanation based on storage conflicts with the facts: a storage-based account predicts that an effect should be immediately observable after the *wh*-phrase and would have to say that storing a less informative representation is more difficult. The other potential candidate explanations, therefore, ultimately lack motivation and require ad hoc or even illogical assumptions.

These results should be balanced against other clear evidence that suggests that the use of shorter and less informative linguistic expressions is often times more efficient and preferable in discourse and writing (Ariel, 1990, 2001). Discourse entities are not perpetually referred to with the same degree of detail. After establishing the identity of individuals and objects in a discourse, shorter and less precise forms are adopted. I suggest that, as with many linguistic processes, there is an on-line competition for using more or less infor-

mative expressions. Shorter and less informative expressions are licensed when their intended referents are highly salient and, hence, easy to retrieve from memory. At other times, longer and more informative expressions are called for to introduce discourse entities, to establish or change discourse topics, and identify the most significant discourse entities. The longer and more distinctive linguistic expressions effectively boost the activation of the associated discourse representations, meaning that future references to those representations can be more economical. These two claims reduce to saying that linguistic form both reflects and attenuates retrievability from memory. Hence, adding an account of memory facilitation to our knowledge about discourse salience creates a more complete picture of how the "accessibility" of discourse entities fluctuates throughout discourse.

The evidence presented here also constitutes at least a potential explanation for the acceptability paradigms raised in linguistic research such as Pesetsky (1987, 2000) and Kluender (1992), where differences in extracted elements create contrasts in acceptability. As processing difficulty can significantly impact acceptability (Chomsky & Miller, 1963; Fanselow & Frisch, 2004), the pattern of more informative fillers increasing the acceptability of certain filler-gap dependencies may partly derive from the decreased processing difficulty, due to easier retrieval. Indeed, other experimental data (Hofmeister et al., 2007) confirm that acceptability and processing difficulty measures converge on the same preference for more informative *wh*-phrases in multiple *wh*-questions, i.e. evidence of easier processing at the retrieval site correlates with higher acceptability ratings.

In sum, the ideas developed here stem from the consideration that some kinds of mental representations are more easily recalled than others. While many factors dynamically contribute to retrievability, I claim that the quantity of information in a linguistic phrase plays a significant role. This research consequently seeks to establish a systematic relationship between how information is encoded and how well that information is remembered later for purposes of natural language comprehension. The implications of this theory extend to any inquiry that examines the choice between more or less complicated locutions and how this affects subsequent language processing.

References

Alexopoulou, T., & Keller, F. (2003). Linguistic complexity, locality, and resumption. In *Proceedings of WCCFL22*. Somerville, MA: Cascadilla Press.

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

Ariel, M. (2001). Accessibility theory: an overview. In T. Sanders, J. Schilperoord, & W. Spooren (Eds.), *Text representation: Linguistic and psycholinguistic aspects*. Amsterdam: John Benjamins.

Bradshaw, G., & Anderson, J. (1982). Elaborative encoding as an explanation of levels of processing. *Journal of Verbal Learning and Verbal Behavior*, *21*, 165-174.

Chomsky, N. (1973). Conditions on transformations. In S. Anderson & P. Kiparsky (Eds.), *A festschrift for Morris Halle* (p. 232-86). New York: Holt, Reinhart & Winston.

Chomsky, N., & Miller, G. (1963). Introduction to the formal analysis of natural languages. In R. Luce, R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology* (p. 269-321). New York: Wiley.

Deane, P. (1991). Limits to attention: a cognitive theory of island phenomena. *Cognitive Linguistics*, *2*(1), 1-63.

Fanselow, G., & Frisch, S. (2004). Effects of processing difficulty on judgments of acceptability. In G. Fanselow, C. Fery, M. Schlesewsky, & R. Vogel (Eds.), *Gradience in grammar*. Oxford: Oxford University Press.

Frazier, L., & Clifton, C. (2002). Processing 'd-linked' phrases. *Journal of Psycholinguistics*, *31*, 633-659.

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

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

Hawkins, J. (1999). Processing complexity and filler-gap dependencies across grammars. *Language*, *75*, 244-285.

Hofmeister, P., Jaeger, T., Sag, I., Arnon, I., & Snider, N. (2007). Locality and accessibility in wh-questions. In *Linguistic evidence: Empirical, computational, and theoretical perspectives*. Berlin: Mouton de Gruyter.

Kluender, R. (1992). Deriving islands constraints from principles of predication. In H. Goodluck & M. Rochemont (Eds.), *Island constraints: Theory, acquisition and processing* (p. 223-258). Dordrecht: Kluwer.

Kluender, R. (1998). On the distinction between strong and weak islands: a processing perspective. In P. Culicover & L. McNally (Eds.), *Syntax and semantics 29: The limits of syntax* (p. 241-279). San Diego, CA: Academic Press.

Pesetsky, D. (1987). Wh-in-situ: movement and unselective binding. In E. Reuland & A. ter Meulen (Eds.), *The representation of (in)definiteness*. Cambridge: MIT Press.

Pesetsky, D. (2000). *Phrasal movement and its kin*. Cambridge: MIT Press.

Ross, J. R. (1967). *Constraints on variables in syntax*. PhD Thesis. MIT.

Van Dyke, J., & McElree, B. (2006). Retrieval interference in sentence comprehension. *Journal of Memory and Language*, *55*, 157-166.

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