Published in final edited form as:

Lang Cogn Neurosci. 2016 January ; 31(1): 73–79. doi:10.1080/23273798.2015.1036089.

Prediction in the Processing of Repair Disfluencies

Matthew W. Lowder and Fernanda Ferreira

Department of Psychology, University of California, Davis

Abstract

Imagine a speaker who says "Turn left, uh I mean..." Before hearing the repair, the listener is likely to anticipate the word "right" based on the context, including the reparandum "left." Thus, even though the reparandum is not intended as part of the utterance, the listener uses it as information to predict the repair. The issue we explore in this article is how prediction operates in disfluency contexts. We begin by describing the Overlay model of disfluency comprehension, which assumes that the listener identifies a reparandum as such only after a repair is encountered which creates a local ungrammaticality. The Overlay model also allows the reparandum to influence subsequent processing, because the reparandum is not deleted from the final representation of the sentence. A somewhat different model can be developed which assumes a more active, anticipatory process for resolving repair disfluencies. On this model, the listener might predict the likely repair when the speaker becomes disfluent, or even identify a reparandum if the word or word string seems inconsistent with the speaker's intention. Our proposal is that the prediction can be made using the same mechanism involved in the processing of contrast, in which a listener uses contrastive prominence to generate likely alternates (the contrast set). We suggest that these two approaches to disfluency processing are not inconsistent: Successful repair processing requires listeners to use statistical and linguistic evidence to identify a reparandum and to integrate the repair, and the lingering of the reparandum is due to the coexistence in working memory of the reparandum, the repair, and unselected members of the contrast set.

A central goal of psycholinguistics is to understand the cognitive mechanisms that enable us to process and comprehend spoken utterances. Pursuit of this goal has traditionally involved presenting listeners with language that has been scripted to be perfectly fluent and error-free, much like what one would expect to hear in a well-rehearsed speech or to read in a carefully proofread text. However, spontaneous speech is far from perfect. Common disfluencies in everyday language include filled pauses (e.g., *uh, um*), repetitions (e.g., *Turn, Turn left*), and repairs (e.g., *Turn left, uh I mean right*). Although several lines of work have shown that listeners' processing and interpretation of language are influenced by disfluencies (e.g., Arnold, Fagnano, & Tanenhaus, 2003; Arnold, Hudson Kam, & Tanenhaus, 2007; Arnold, Tanenhaus, Altmann, & Fagnano, 2004; Bailey & Ferreira, 2003; Barr & Seyfeddinipur, 2010; Brennan & Schober, 2001; Ferreira, Lau, & Bailey, 2004; Kidd, White, & Aslin, 2011; Lau & Ferreira, 2005; Maxfield, Lyon, & Silliman, 2009), an underexplored area involves understanding the role prediction plays during the processing of disfluent speech,

particularly in the case of repair disfluencies. For example, it seems likely that upon hearing an utterance such as, “*Turn left, uh I mean...*”, the listener will use the reparandum (i.e., *left*) to predict the repair (i.e., *right*) before it is spoken, even though the reparandum was not intended to be part of the utterance. In this paper, we discuss different theoretical approaches to the processing of disfluencies and propose how these approaches might be integrated into a coherent framework to explain how prediction operates in contexts involving various types of disfluencies.

Overlay Model of Disfluency Processing

Consider another example of a repair disfluency, illustrated in (1a). The parser’s initial attempt to build a coherent structure is interrupted by the local ungrammaticality of the input (“*you should put—you should drop*”). One might assume that this ungrammaticality would serve as a cue to the parser to completely eliminate the initial verb *put* from the structure and replace it with the verb *drop* so that only the intended meaning of the revised utterance remains. However, this is not exactly what happens. Note that the verbs in this example contain different argument structures; that is, whereas *put* requires three arguments (i.e., subject, object, goal), the verb *drop* requires only two (i.e., subject and object) and can optionally take a third argument (i.e., goal). Ferreira et al. (2004) demonstrated that subtle differences in the argument structure of a reparandum and a repair can have important consequences for how the entire utterance is interpreted. Specifically, they showed that disfluent utterances where the verb contained in the reparandum had a more complex argument structure than the verb in the repair (1a) were judged to be less acceptable than fluent utterances containing only the verb with the simpler argument structure (e.g., *drop*). In contrast, disfluent utterances where the verb contained in the reparandum had a less complex argument structure than the verb in the repair (1b) were judged to be *more* acceptable than fluent utterances containing the verb with the more complex argument structure (e.g., *put*). What this suggests is that the reparandum is not completely erased from the parse, but rather traces of the reparandum linger and continue to influence the listener’s ultimate interpretation of the sentence.

(1a) John says you should put—you should drop the ball.

(1b) John says you should drop—you should put the ball.

Additional support for the notion that properties of the reparandum influence processing comes from Lau and Ferreira (2005), who investigated the effects of repair disfluencies in listeners’ comprehension of garden-path sentences, as in (2), where the target verb (e.g., *selected*) was ambiguous between a main-verb or reduced-relative interpretation. The critical manipulation was that the verb in the reparandum either unambiguously signaled the correct reduced-relative interpretation of the target verb (2a) or was itself ambiguous between the two interpretations (2b). Results showed that listeners more easily understood sentences like (2a) than sentences like (2b), suggesting that the passive structure activated by the reparandum in (2a) lingered even after it was replaced by a new verb and that this information was used by the parser to resolve the ambiguity (see also Bailey & Ferreira, 2003, for evidence that listeners use filled pauses in their processing of garden-path sentences).

(2a) The little girl chosen, uh, selected for the role celebrated with her friends.

(2b) The little girl picked, uh, selected for the role celebrated with her friends.

These findings led to the development of the Overlay model of disfluency processing (Ferreira & Bailey, 2004; Ferreira et al., 2004), which operates within a Lexicalized Tree-Adjoining Grammar (LTAG) framework (Joshi & Schabes, 1997). LTAG assumes that as the comprehender retrieves the lexical entry of each word during sentence processing, it also retrieves an elementary syntactic tree anchored to that word. As processing proceeds incrementally, these elementary trees are combined through operations called Substitution and Adjunction. The Overlay model proposes that upon encountering a local ungrammaticality triggered by a repair disfluency (e.g., *chosen, uh, selected*), syntactic trees for both the reparandum and repair are retrieved, and when the parser cannot combine them at an appropriate substitution site, they are combined via an operation called Overlay in which the repair tree is superimposed over the reparandum tree. When the reparandum and repair have different argument structures, there is imperfect structural overlap, which allows traces of the reparandum tree to linger and continue to influence the comprehender's interpretation of the sentence and offline judgments of acceptability. In addition to the structural differences between the reparandum and repair, there may also be lexical differences between the words anchoring the reparandum and repair trees, and these could also potentially influence processing. The stimuli used in the experiments were chosen to have similar lexical meanings (e.g., *chosen/selected*) in order to isolate structural influences, but to the extent that the meanings of the lexical anchors differ, the Overlay model assumes that those semantic properties will linger as well. This prediction has not been directly tested yet, however, and so it is not known whether the meanings of the reparandum and repair would merge or whether sometimes the reparandum would fail to be replaced in the final representation (leading to memory errors, for example).

A major assumption of the Overlay model is that the process of superimposing the two structures and attempting to derive the correct meaning does not begin until a local ungrammaticality is detected. In this way, the mechanisms that contribute to comprehension of disfluencies have been likened to those that contribute to the comprehension of garden-path sentences, in the sense that both contain an ambiguous region that is potentially misinterpreted, followed by a point of disambiguation that triggers a process of reanalysis. Work that has carefully probed comprehenders' interpretations of garden-path sentences has revealed that the initial interpretation often has lingering effects even after the disambiguating information has been encountered (e.g., Christianson, Hollingworth, Halliwell, & Ferreira, 2001; Ferreira, Christianson, & Hollingworth, 2001), leading to linguistic representations that are incomplete, inaccurate, or otherwise simply "good enough" (Ferreira, Bailey, & Ferraro, 2002; Ferreira & Patson, 2007). The Overlay model assumes that the processing of repair disfluencies proceeds in a similar fashion, such that the listener adopts an initial interpretation of the ambiguous region (i.e., the segment that will ultimately be the reparandum), encounters an explicit self-correction (e.g., *uh, I mean*) or other disruption (e.g., *chosen selected*), which serves as a cue that the initial interpretation was incorrect and must be reanalyzed; however, as we have seen, traces of the reparandum (e.g., its argument structure) can linger and continue to influence processing. Importantly,

this characterization of the processing of disfluencies can be thought of as a “backward-looking” model in the sense that reanalysis and reinterpretation of an utterance begin only after an overt cue in the input signals that a problem has been encountered and must be resolved.

Bayesian/Noisy Channel Models of Language Processing

A very different theoretical approach to the processing of disfluencies can be developed from recent work that describes language comprehension as the transmission of a linguistic signal across a noisy channel (Gibson, Bergen, & Piantadosi, 2013; Gibson, Piantadosi, Brink, Bergen, Lim, & Saxe, 2013). “Noise” in a Noisy-Channel framework refers to any of the distortions inherent to natural language that might affect the comprehender’s ability to process it, such as typos and poor handwriting in the domain of written language, or background noise, nonnative speech, and producer errors in the domain of spoken language. A Noisy-Channel framework thus assumes that linguistic input is imperfect and that comprehenders are “forward-looking” in the sense that they predict and anticipate the speaker’s meaning, combining the actual linguistic input with expectations about what sorts of input seem most plausible given the context, speaker characteristics, and other factors. Further, a Noisy-Channel account proposes that comprehenders possess a mechanism that actively corrects errors, especially if the wording of a more plausible interpretation is very close to the wording of the perceived input. For example, if the listener hears an utterance that is deemed implausible (e.g., *The mother gave the candle the daughter*), he or she will likely assume that the speaker misspoke and meant to say *The mother gave the candle TO the daughter*, or the comprehender may assume that the speaker spoke correctly but that perhaps the word “to” was obscured by noise in the environment (Gibson, Bergen, & Piantadosi, 2013).

The predictions of the Noisy-Channel model easily extend to the processing of repair disfluencies. For example, consider the utterance, “*Put the milk in the stove, uh, the fridge please.*” Under a Noisy-Channel framework, the listener is constantly weighing the perceived input against the plausibility of that input and assumptions about the speaker’s meaning. In this example, the much higher probability that the speaker would want the milk to go in the refrigerator instead of the stove should lead the listener to anticipate that this is what the speaker means and make a mental correction to the speech signal even before the speaker articulates the revision. Under this account, when the speaker does explicitly make the correction, the listener’s task is to determine whether the anticipated repair and the actual repair match and to engage in a process of accommodation only if they do not. Thus, the Noisy-Channel approach to disfluencies assumes that listeners probabilistically evaluate incoming speech against their expectations, rapidly detect any input that is likely to be a reparandum, and actively predict the possible repair. This process can be influenced by any contextual information available to the listener, including environmental cues. For example, it seems that the listener would be even more likely to identify “stove” as a reparandum in the above example if the speaker is looking at the refrigerator while speaking.

The Role of Prediction in the Processing of Disfluencies

Up to now, most research examining the effects of disfluency on incremental comprehension processes has focused on filled pauses such as *um* and *er*. This work has shown that listeners who hear sentences while at the same time viewing images on a computer display make anticipatory eye movements to previously unmentioned objects (Arnold et al., 2004). They also anticipate mentions to objects that are difficult to label or describe (Arnold et al., 2007). The logic is that listeners know that speakers tend to produce filled pauses when they have difficulty with production, and listeners are able to use this correlation to predict the upcoming word or phrase. This ability seems to emerge quite early in development: Children as young as two years of age attend to filled pauses in speech and use them as cues to predict speakers' intended referents (Kidd et al., 2011). Evidence also suggests that listeners use filled pauses to make syntactic predictions (Bailey & Ferreira, 2004, 2007). These studies start from the observation that speakers tend to be disfluent before heavier and more complex syntactic constituents (e.g., Clark & Wasow, 1998; Ford, 1982; Hawkins, 1971; Shriberg, 1996). This correlation would then allow listeners to correctly choose among alternative syntactic analyses for an ambiguous phrase, as in *Put the apple on the towel in the box*. Bailey and Ferreira (2007) found that filled pauses before *apple* led listeners to treat *on the towel* as a modifier of *apple*, but filled pauses after *towel* led listeners to treat *in the box* as a modifier of *towel* (i.e., [the apple on the towel][in the box] versus [the apple] [on the towel in the box]), as reflected in anticipatory eye movements toward a set of objects consisting of a simple apple, an apple on a towel, a towel, and a towel in a box.

It seems plausible that prediction would also play a role in the resolution of repair disfluencies, but for these, the mechanism may be somewhat different. Recall that the challenge for the listener is to recognize a reparandum as such, and to replace the content of the reparandum with the speaker's intended repair. As discussed in the section describing the Overlay model, one approach to solving this problem is to work backwards from the repair to the reparandum, using syntactic and lexical constraints to identify the domain of each. A different approach would be for the listener to identify a reparandum based on the presence of an improbable word or string (e.g., *Put the milk in the stove...*) or because the speaker signals that an error has been made (e.g., *Please hand me the milk, uh I mean...*) and then use this information to anticipate a repair before it is spoken. In the next section, we discuss how a Noisy-Channel model of language processing can be combined with our understanding of the processing of focus to suggest a possible mechanism for generating these predictions.

Noisy-Channel Models, Focus, and Repair Disfluencies

Recall that Noisy-Channel models assume listeners engage in an ongoing process of assessing the likelihood of the input and mentally repairing words and phrases that seem inconsistent with the speaker's communicative intention. Applying these models to repair disfluencies, we propose that when listeners identify a reparandum, they use semantic information about this word, along with relevant contextual information, to predict a set of likely repairs. More specifically, our claim is that the reparandum cues the listener to generate a set of contrastive alternatives that seem like plausible candidates for the repair. A

variety of linguistic and psycholinguistic perspectives have suggested that language comprehenders are sensitive to cues that mark linguistic focus, such as syntactic structure, prosody, or prior discourse context (e.g., Chomsky, 1971; Haliday, 1967; Jackendoff, 1972; Rooth, 1992), and can then use this information to make inferences or otherwise facilitate processing (e.g., Birch & Garnsey, 1995; Cutler & Fodor, 1976; Fraundorf, Watson, & Benjamin, 2010; Gergely, 1992; Hornby, 1974; Sedivy, 2002). For example, if a speaker says, “*I don't want the BLUE pen...*” (particularly with prosodic emphasis on *blue*), the listener identifies *blue* as focused and generates a set of alternates that are likely to be mentioned—red, green, black, etc.—allowing the listener to more easily process the utterance describing what type of pen the speaker does want (e.g., “*I want the red pen instead*”).

Now consider a speaker who says, “*I want the blue pen, uh I mean the...*” Our claim is that listeners engage in a process very similar to the one just described: When *blue* is identified as a reparandum, listeners generate an alternate set containing a list of likely colors for pens. Similarly, if a speaker says, “*I just ate the apple, uh I mean...*”, we propose that the listener interprets this utterance to mean that the speaker did not eat the apple, and then the listener uses this information to predict likely repairs that stand in semantic contrast to the reparandum (e.g., orange, banana, pear, kiwi, etc.). Importantly, certain items in this set may be more highly activated than others. A particular item may be highly activated because it represents a strong semantic associate of the reparandum (e.g., orange). Alternatively, in some cases the listener may have relevant knowledge about the speaker that can be used to predict the repair (e.g., the listener knows that the speaker often eats kiwis).

The above example illustrates how the Noisy-Channel model can be combined with this focus mechanism to explain the use of prediction in the processing of repair disfluencies. That is, once the listener receives some indication that the linguistic input contains an error, he or she uses semantic information about the reparandum, along with any relevant contextual information to generate a set of likely repairs. Similarly, we propose that a listener who has identified a sequence as low probability has also identified a potential reparandum, and at that point activates an alternate set, before any explicit evidence becomes available indicating that the speaker has been disfluent. In this context, consider once again our example “*Put the milk in the stove.*” Given the very low probability that the speaker intended to say “*stove*” in this context, the listener would mentally correct “*stove*” to “*fridge*” or “*refrigerator*,” and if the speaker corrects him- or herself, the listener's task is to reconcile the explicit correction with the one generated covertly. Crucially, our proposal is that these mental corrections are based on the creation of an alternate set similar to alternate sets generated during the processing of semantic focus.

As mentioned above, a large body of psycholinguistic literature has shown that language comprehenders are sensitive to manipulations of focus signaled via syntax (e.g., Birch & Rayner, 1997, 2010; Lowder & Gordon, under review; Morris & Folk, 1998), prosody (e.g., Arnold, 2008; Birch & Clifton, 1995, 2002; Cutler & Foss, 1977; Dahan, Tanenhaus, & Chambers, 2002), or discourse context (e.g., Benatar & Clifton, 2014; Cutler & Fodor, 1979; Sturt, Sanford, Stewart, & Dawydiak, 2004). In addition, recent work using the visual-world paradigm (Tanenhaus, Spivey-Knowlton, Eberhardt, & Sedivy, 1995) has shown that

negation operators (e.g., *not*) function as effective focus cues, signaling listeners to rapidly shift their focus from the negated entity to an alternate entity when a strongly activated alternate is available (Orenes, Beltrán, & Santamaría, 2014). A question for future work involves understanding the extent to which repair disfluencies are effective focusing devices that cue listeners to generate predictions. Data from the visual-world paradigm could be particularly useful in answering this question, as this method has proven to be very sensitive for measuring the predictions generated by listeners during online language processing (e.g., Altmann & Kamide, 1999; Corley, 2010; Kamide, Altmann, & Haywood, 2003; Kaiser & Trueswell, 2004). An approach of this sort could analyze listeners' fixation patterns to items in a visual array while they hear disfluent utterances (e.g., *At lunch, the boy ate the apple, uh I mean...*) as well as utterances that contain negation (e.g., *At lunch, the boy ate not the apple, but rather...*). Fixation patterns to images corresponding with the named item (e.g., *apple*), a closely related alternate (e.g., *orange*), and a less common alternate (e.g., *kiwi*) could be compared to examine at what point during processing listeners begin generating predictions. An important possible difference between these two types of constructions is that the negation operator clearly marks the entity that is meant to stand in contrast to the alternate set (e.g., *not the apple*), whereas the editing interval in the disfluent utterance (e.g., *uh, I mean*) does not explicitly signal what the reparandum will be. Although it seems likely in this example that "*apple*" will be the reparandum (e.g., *At lunch, the boy ate the apple, uh I mean, the orange*), different reparanda are also possible (e.g., *At lunch, the boy ate the apple, uh I mean, the girl ate the apple*). Thus, an important goal of future work will be to understand not only what information the listener uses to generate predictions about the upcoming repair, but also what information the listener uses to anticipate which portion of the sentence will be the site of the reparandum.

It should also be noted that reparandum-repair disfluencies do not always indicate that the speaker has made an error, but instead sometimes indicate the speaker's desire to resolve potential ambiguity or otherwise use a more appropriate word (E-Repairs versus A-Repairs in Levelt's, 1983, terminology). For example, consider a speaker who says "*I just ate the apple, uh I mean the Granny Smith apple.*" In this case, the speaker has decided to replace the original term with a more specific one, perhaps to highlight this particular apple from another set of apples. It is unclear at this point how mechanisms of prediction might operate in contexts such as this. On the one hand, it may be the case that listeners include possible A-Repairs of the reparandum in the alternate set, along with possible E-Repairs. Instead, it may be the case that the alternate set typically only contains E-Repairs to allow for more rapid and efficient prediction. In this latter case, there may be a processing cost when the reparandum is replaced with an A-Repair rather than an E-Repair. We acknowledge that these possibilities are highly speculative, and we leave it to future research to explore the extent to which mechanisms of prediction in these two types of repair disfluencies are similar or different from one another.

Conclusion

As we have argued, the Overlay and Noisy-Channel models differ with respect to their basic assumptions about the mechanisms that allow listeners to identify and resolve disfluencies. The Noisy-Channel model captures the ability of the comprehension system to identify a

word or word string as unlikely given the speaker's communicative intentions, allowing the listener to anticipate a repair or to mentally correct a speech error. The Noisy-Channel model also accounts for listeners' tendencies to treat input probabilistically; perceptual and production systems are error-prone, and the channel over which communication takes place may mask or distort the linguistic signal. In contrast, the Overlay model is designed to explain how the comprehension system deals with the reparandum once a repair has been identified and processed: Rather than being erased, the reparandum is covered up by the repair, which allows the properties of the reparandum to leak through and affect the listener's comprehension of the utterance.

We hope it is clear, then, that these two approaches are not mutually exclusive, but rather both may explain important aspects of disfluency processing. Rapid and efficient language comprehension requires listeners to use linguistic and contextual information to generate predictions about upcoming material. Listeners may be particularly likely to generate predictions during the processing of repair disfluencies, where there is often a relatively long period of time between the reparandum and the repair, during which the speaker attempts to retrieve the intended word and the listener uses information about the reparandum, the speaker, and the context to uncover the speaker's meaning and derive a set of possible repairs that the speaker believes are consistent with that meaning. When the speaker continues with the utterance, the listener will attempt to match the actual repair with the predicted repair, making adjustments as necessary. Importantly, however, we have seen that the reparandum is not completely eliminated from the listener's representation of the utterance, and information about the reparandum can continue to influence the listener's ultimate interpretation of the sentence. This persistence of the reparandum may in part be attributable to its prominent role in generating an alternate set; that is, the reparandum lingers because not only is it not filtered or ignored, but it is used to facilitate processing of the repair via the creation of the alternate set. These items then coexist in working memory during online sentence comprehension: the reparandum, the repair, and the alternate set. Given that the auditory signal is ephemeral and that perceptual systems and working memory operate under a range of constraints, it is not surprising that listeners will sometimes retain a representation of the reparandum in their final interpretation of an utterance containing a repair disfluency. In future work, our goal is to discover direct evidence for the mechanisms we have proposed here, and to develop a specific model of how the reparandum is processed and inhibited during incremental language processing.

Acknowledgments

This work was supported by R56 DC013545 from NIDCD.

References

- Altmann GTM, Kamide Y. Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition*. 1999; 73:247–264. [PubMed: 10585516]
- Arnold JE. *THE BACON* not *the bacon*: How children and adults understand accented and unaccented noun phrases. *Cognition*. 2008; 108:69–99. [PubMed: 18358460]
- Arnold JE, Fagnano M, Tanenhaus MK. Disfluencies signal theee, um, new information. *Journal of Psycholinguistic Research*. 2003; 32:25–36. [PubMed: 12647561]

- Arnold JE, Hudson Kam CL, Tanenhaus MK. If you say *thee uh* you are describing something hard: The on-line attribution of disfluency during reference comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2007; 33:914–930.
- Arnold JE, Tanenhaus MK, Altmann R, Fagnano M. The old and thee, uh, new. *Psychological Science*. 2004; 15:578–581. [PubMed: 15327627]
- Bailey KGD, Ferreira F. Disfluencies affect the parsing of garden-path sentences. *Journal of Memory & Language*. 2003; 49:183–200.
- Bailey, KGD.; Ferreira, F. The processing of filled pause disfluencies in the visual world. In: van Gompel, RPG.; Fischer, MH.; Murray, WS.; Hill, RL., editors. *Eye movements: A window on mind and brain*. Oxford, UK: Elsevier; 2007. p. 485-500.
- Barr DJ, Seyfeddinipur M. The role of fillers in listener attributions for speaker disfluency. *Language and Cognitive Processes*. 2010; 25:441–455.
- Benatar A, Clifton C Jr. Newness, givenness and discourse updating: Evidence from eye movements. *Journal of Memory and Language*. 2014; 71:1–16.
- Birch S, Clifton C Jr. Focus, accent, and argument structure: Effects on language comprehension. *Language and Speech*. 1995; 38:365–391. [PubMed: 8816086]
- Birch S, Clifton C Jr. Effects of varying focus and accenting adjuncts on the comprehension of utterances. *Journal of Memory and Language*. 2002; 47:571–588.
- Birch SL, Garnsey SM. The effect of focus on memory for words in sentences. *Journal of Memory and Language*. 1995; 34:232–267.
- Birch S, Rayner K. Linguistic focus affects eye movements during reading. *Memory & Cognition*. 1997; 25:653–660. [PubMed: 9337583]
- Birch S, Rayner K. Effects of syntactic prominence on eye movements during reading. *Memory & Cognition*. 2010; 38:740–752. [PubMed: 20852237]
- Brennan SE, Schober MF. How listeners compensate for disfluencies in spontaneous speech. *Journal of Memory and Language*. 2001; 44:274–296.
- Christianson K, Hollingworth A, Halliwell JF, Ferreira F. Thematic roles assigned along the garden path linger. *Cognitive Psychology*. 2001; 42:368–407. [PubMed: 11368528]
- Chomsky, N. Deep structure, surface structure, and semantic interpretation. In: Steinberg, DD.; Jakobovits, LA., editors. *Semantics: An interdisciplinary reader in philosophy, linguistics, and psychology*. New York: Cambridge University Press; 1971. p. 183-216.
- Clark HH, Wasow T. Repeating words in spontaneous speech. *Cognitive Psychology*. 1998; 37:201–242. [PubMed: 9892548]
- Corley M. Making predictions from speech with repairs: Evidence from eye movements. *Language and Cognitive Processes*. 2010; 25:706–727.
- Cutler A, Fodor JA. Semantic focus and sentence comprehension. *Cognition*. 1979; 7:49–59. [PubMed: 436402]
- Cutler A, Foss DJ. On the role of sentence stress in sentence processing. *Language and Speech*. 1977; 20:1–10. [PubMed: 592948]
- Dahan D, Tanenhaus MK, Chambers CG. Accent and reference resolution in spoken-language comprehension. *Journal of Memory and Language*. 2002; 47:292–314.
- Ferreira F, Bailey KGD. Disfluencies and human language comprehension. *Trends in Cognitive Sciences*. 2004; 8:231–237. [PubMed: 15120682]
- Ferreira F, Bailey KGD, Ferraro V. Good-enough representations in language comprehension. *Current Directions in Psychological Science*. 2002; 11:11–15.
- Ferreira F, Christianson K, Hollingworth A. Misinterpretations of garden-path sentences: Implications for models of sentence processing and reanalysis. *Journal of Psycholinguistic Research*. 2001; 30:3–20. [PubMed: 11291182]
- Ferreira F, Lau EF, Bailey KGD. Disfluencies, language comprehension, and tree adjoining grammars. *Cognitive Science*. 2004; 28:721–749.
- Ferreira F, Patson N. The good enough approach to language comprehension. *Language and Linguistics Compass*. 2007; 1:71–83.

- Ford, M. Sentence planning units: Implications for the speaker's representation of meaningful relations underlying sentences. In: Bresnan, J., editor. *The mental representation of grammatical relations*. Cambridge, MA: MIT Press; 1982. p. 798-827.
- Fraundorf SH, Watson DG, Benjamin AS. Recognition memory reveals just how CONTRASTIVE contrastive accenting really is. *Journal of Memory and Language*. 2010; 63:367–386. [PubMed: 20835405]
- Gergely, G. Focus-based inferences in sentence comprehension. In: Sag, IA.; Szabolcsi, A., editors. *Lexical matters*. Stanford, CA: Center for the Study of Language and Information; 1992. p. 209-240.
- Gibson E, Bergen L, Piantadosi ST. Rational integration of noisy evidence and prior semantic expectations in sentence interpretation. *Proceedings of the National Academy of Sciences*. 2013; 110:8051–8056.
- Gibson E, Piantadosi ST, Brink K, Bergen L, Lim E, Saxe R. A noisy-channel account of crosslinguistic word-order variation. *Psychological Science*. 2013; 24:1079–1088. [PubMed: 23649563]
- Halliday MAK. Notes on transitivity and theme in English: Part 2. *Journal of Linguistics*. 1967; 3:199–244.
- Hawkins PR. The syntactic location of hesitation pauses. *Language and Speech*. 1971; 14:277–288. [PubMed: 5097643]
- Hornby PA. Surface structure and presupposition. *Journal of Verbal Learning and Verbal Behavior*. 1974; 13:530–538.
- Jackendoff, RS. *Semantic interpretation in generative grammar*. Cambridge, MA: MIT Press; 1972.
- Joshi, AK.; Schabes, Y. Tree-adjoining grammars. In: Rozenberg, G.; Salomaa, A., editors. *Handbook of formal languages*. Berlin: Springer-Verlag; 1997. p. 69-123.
- Kaiser E, Trueswell JC. The role of discourse context in the processing of a flexible word-order language. *Cognition*. 2004; 94:113–147. [PubMed: 15582623]
- Kamide Y, Altmann GTM, Haywood SL. The time course of prediction in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language*. 2003; 49:133–156.
- Kidd C, White KS, Aslin RN. Toddlers use speech disfluencies to predict speakers' referential intentions. *Developmental Science*. 2011; 14:925–934. [PubMed: 21676111]
- Lau EF, Ferreira F. Lingering effects of disfluent material on comprehension of garden path sentences. *Language and Cognitive Processes*. 2005; 20:633–666.
- Levelt WJM. Monitoring and self-repair in speech. *Cognition*. 1983; 14:41–104. [PubMed: 6685011]
- Lowder MW, Gordon PC. Focus takes time: Structural effects on reading. (under review). *Manuscript under review*.
- Maxfield ND, Lyon JM, Silliman ER. Disfluencies along the garden path: Brain electrophysiological evidence of disrupted sentence processing. *Brain & Language*. 2009; 111:86–100. [PubMed: 19765813]
- Morris RK, Folk JR. Focus as a contextual priming mechanism in reading. *Memory & Cognition*. 1998; 26:1313–1322. [PubMed: 9847554]
- Orenes I, Beltrán D, Santamaría C. How negation is understood: Evidence from the visual world paradigm. *Journal of Memory and Language*. 2014; 74:36–45.
- Rooth M. A theory of focus interpretation. *Natural Language Semantics*. 1992; 1:75–116.
- Sedivy JC. Invoking discourse-based contrast sets and resolving syntactic ambiguities. *Journal of Memory and Language*. 2002; 46:341–370.
- Shriberg, EE. *Proceedings of International Conference on Spoken Language Processing, Addendum*. Philadelphia, PA: 1996. Disfluencies in SWITCHBOARD; p. 11-14.
- Sturt P, Sanford AJ, Stewart A, Dawydiak E. Linguistic focus and good-enough representations: An application of the change-detection paradigm. *Psychonomic Bulletin & Review*. 2004; 11:882–888. [PubMed: 15732698]

Tanenhaus MK, Spivey-Knowlton MJ, Eberhard KM, Sedivy JC. Integration of visual and linguistic information in spoken language comprehension. *Science*. 1995; 268:1632–1634. [PubMed: 7777863]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript