Language and Linguistics Compass 1/1–2 (2007): 71–83, 10.1111/j.1749-818x.2007.00007.x

The 'Good Enough' Approach to Language Comprehension

Fernanda Ferreira* and Nikole D. Patson
University of Edinburgh

Abstract

Ferreira and colleagues argued that the language comprehension system creates syntactic and semantic representations that are merely 'good enough' (GE) given the task that the comprehender needs to perform. GE representations contrast with ones that are detailed, complete, and accurate with respect to the input. In this article, we review the original argument for GE processing, and we present new evidence that supports the concept: first, local interpretations are computed, which can interfere with global ones; second, new findings based on the recording of event-related potentials show the use of simple heuristics rather than compositional algorithms for constructing sentence meaning; and recent studies show that the comprehension system has mechanisms for handling disfluencies, but they work imperfectly. We argue that the GE approach to language comprehension is similar to the use of fast and frugal heuristics for decision-making, and that future research should explore this connection more thoroughly.

In 2002, Ferreira and colleagues suggested that the representations created during language comprehension were just 'good enough' (GE) (Ferreira et al. 2002). The fundamental idea was to challenge the assumption that the overarching goal of the language comprehension system is to deliver an accurate and detailed representation of the speaker's utterance. Ferreira et al. argued that instead the system's responsibility is to create representations that are suitable for the task that the listener wants to perform with the help of the linguistic input. Typically, the listener's task is to provide an appropriate followup to the speaker's contribution (i.e., to maintain a dialogue), but of course there are other possibilities – for example, the listener might be required to execute a motor action, or merely nod politely at appropriate moments. Their point was that listeners are rarely required to prove the accuracy or detailed nature of their understanding of some utterance, and this might have important implications for the architecture of the language processing system.

To be more specific, consider a sentence such as *The singer saw the audience member with the binoculars*. This sentence is ambiguous; the binoculars are either in the singer's or in the audience member's possession. A great deal of psycholinguistic work has been devoted to trying to determine how these

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

sorts of syntactic ambiguities are resolved. Both in models of comprehension that assume syntactic priority (Frazier & Fodor 1978; Frazier 1979; Frazier & Rayner 1982) and in those that allow unlimited interaction among knowledge sources with no special status for syntax (Trueswell et al. 1993; MacDonald et al. 1994; Stevenson 1994; Spivey & Tanenhaus 1998), it is assumed that eventually all relevant sources of information will be consulted and used to arrive at a definitive resolution of this and similar ambiguities. In the example, because it is common knowledge that singers usually have no need for binoculars, whereas audience members might use them to see the stage clearly, the ambiguity should be resolved so that the phrase *with the binoculars* attaches low in the structure, thus syntactically indicating the correct modification relationship. But in their proposal for 'GE processing', Ferreira et al. suggested that the system might not resolve the ambiguity at all, or it might arrive at an interpretation that is inconsistent with some relevant sources of knowledge. They also claimed that psycholinguistic experiments might bias comprehenders to work harder than usual to arrive at a single and definitive interpretation, either because participants are typically asked comprehension questions that reveal the nature of the interpretation, or simply because of the demand characteristics of the experimental situation ('why would I be asked to read these sentences if the experimenter didn't want me to try to understand what they mean?'). But in the normal circumstances of communication, it might be that nothing turns on resolving this ambiguity. Moreover, utterances in dialogue come quickly one after another, and so the system might not have time to consider all relevant sources of information and compute a specific and detailed structure for each one. The system might instead rely on 'GE representations' that are refined if necessary.

In this article, we will elaborate on the basic concept. First, we will briefly summarize the evidence for GE processing in language comprehension. Second, we will describe recent studies that provide new evidence for GE comprehension. These studies address phenomena such as shallow and local processing, the detection of anomaly and ungrammaticality, and the comprehension of utterances containing disfluencies. We end by discussing an intriguing future direction for the idea of GE comprehension – we suggest that it should be connected to work in decision-making that also challenges the idea of 'unbounded rationality' and assumes instead that the cognitive system relies on a small set of fast and frugal heuristics to accomplish its information processing task. We argue that a model of parsing that assumes a minimal effort principle for initial parsing is most compatible with this approach.

Initial Evidence for GE Comprehension

In 2002, Ferreira et al. argued that two important findings undermine the idea that comprehenders always create accurate and detailed representations. The first was that people often obtain a shallow understanding of a sentence's meaning; the other was that people sometimes outright misunderstand

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

sentences. Because the evidence was reviewed in the 2002 paper and elsewhere, we will summarize it just briefly here.

Imagine that you're asked, 'How many of each type of animal did Moses take on the ark?' If you are like most people, you will respond by saying 'two' rather than by pointing out that the presupposition behind the question is false, because it was Noah who saved the animals from the great flood, not Moses (Erickson & Matteson 1981). Similarly, most people overlook the anomaly in the question 'Where should the authorities bury the survivors?' (survivors of course should not be buried, as they are alive). Thus, it appears that people's comprehension of sentences can be quite shallow. In the case of the Moses illusion, people understand the sentence well enough to realize they are being asked about the biblical figure who played a role in the story about the ark, but not deeply enough to appreciate the difference between Noah and Moses. A similar story holds for the bury-the-survivor example (Barton & Sanford 1993). A variety of studies provides further evidence for incomplete and shallow representations of sentence meaning. Words such as *convent* have both a literal and a figurative use (the building and the institution, respectively), but even when one is more frequent than the other, people have little trouble integrating either meaning when later disambiguating information arrives (Pickering & Frisson 2001). Frazier et al. (1999) found evidence that people do not commit to either a distributed or non-distributed reading of a sentence such as *Mary and John saved $100* (each might have saved $100, or they might have saved $100 together), but instead seem to specify the information later if subsequent text demands a particular interpretation. Sanford and Sturt (2002) summarize much of the evidence for incomplete representations as well as the computational advantages of a system that under specifies initially and fills in information if the details become relevant.

Ferreira and colleagues discovered a type of GE comprehension when they decided to assess people's understanding of so-called garden-path sentences (Christianson et al. 2001; Ferreira et al. 2001). An example of this type of sentence is *While Mary bathed the baby played in the crib*. This sentence is difficult to understand because it is almost always initially mis-analyzed. The phrase *the baby* appears to be the object of *bathed*, but it is in fact the subject of *played*. Most work done up to the point when we did our study assessed reading times for garden-path sentences, but did not look carefully at whether they were ultimately understood correctly. This work revealed that most often they are not: After reading a garden-path sentence like our example, most people will reply 'yes' when asked 'Did Mary bathe the baby?' The experiment ruled out a variety of alternative explanations, including the possibility that what we termed a misinterpretation was simply a likely inference. But one possible concern is that the form of the question might have reinstated the original misparse and the corresponding misunderstanding of the sentence, leading to the high percentage of incorrect 'yes' responses. In work currently in progress (Patson et al. 2006) participants were given

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

similar sentences and then asked to recall them either immediately or after a delay, rather than to reply to yes/no questions. The idea was to avoid contaminating the memory representation with a memory probe that might bias the compehender to reinstate the original misparse of the sentence. The results were remarkably similar to what Christianson et al. reported: Most people recall the sentence with the proposition that Mary bathed the baby. Thus, this phenomenon appears to be robust and is not just a matter of the type of task that is used to assess sentence memory.[1]

Not only do people misunderstand garden-path sentences, it also appears that they fail to get the meaning of standard, unambiguous noncanonical sentences as well, including passives. Ferreira (2003) had college students listen to sentences such as *the dog was bitten by the man*. Afterwards, they were asked to name the agent of the action. Surprisingly, a large proportion of the participants gave the wrong answer, even though the sentence was presented clearly and there was little time pressure to respond. The effect was not attributable to the frequency of the structural form itself, because *it was the man who bit the dog* was understood perfectly, even though the cleft form is much less common than the passive. Instead, it appears that comprehenders of English have a heuristic for understanding basic declarative sentences: They assume that the first noun phrase (NP) is the agent of the action, and the next NP is the entity affected by the action (Bever 1970; Townsend & Bever 2001). This strategy is so compelling, it leads to misunderstandings of quite simple sentences, even in a laboratory situation in which the stimulus is unambiguous and there is ample time in which to process the sentence's structure and meaning. It thus appears that although the comprehension system clearly makes use of syntactic algorithms, it also uses heuristics which generally operate more quickly (Townsend & Bever 2001). The heuristic interpretation might then be selected because it becomes available more quickly, or because the system has some reason for preferring it (e.g., it conforms to real world knowledge).

What's New on the GE Front?

A substantial amount of work published since 2002 provides additional support for the notion of GE comprehension. We will briefly review this research here. First, we consider evidence that people compute local interpretations that are sometimes inconsistent with the overall sentence structure, indicating that the comprehension system tries to construct interpretations over small numbers of adjacent words whenever possible and can be lazy about computing a more global structure and meaning. Second, we discuss new studies using event-related potentials (ERP), which provide more ammunition for the claim that the comprehension system relies on heuristics. Third, we report on work showing that people do not expect spoken utterances to be error free, and in fact seem to have mechanisms for dealing with disfluencies – mechanisms which nonetheless do not work

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

flawlessly. Because of space limitations, we will only be able to provide information about the broad outlines of these investigations. Our goal will be not to evaluate these contributions in detail, but instead to indicate how a number of trends in psycholinguistics are consistent with the GE idea.

First, then, the effects of local coherence (Tabor et al. 2004). Consider the sentence *The coach smiled at the player tossed a frisbee by the opposing team*. The correct, global meaning is that the player was the recipient of the frisbee; however, embedded in the string is the sequence *the player tossed a frisbee*, in which the player is the agent of the frisbee toss. Tabor et al. found that people slowed down when they encountered *tossed*, which Tabor et al. interpret as evidence that the local meaning disrupts processing of the global interpretation. Another way of describing the result is in terms of GE comprehension. The active, main clause structure conforms to the heuristic Bever (1970) and Ferreira (2003) identified, and the meaning that heuristic generates challenges the parser when it tries to integrate the other words in the sentence. Tabor et al. considered this idea in their paper but then rejected it, based on their worry that the Christianson et al. (2001) results might have been due to the probe question. However, the study we mentioned earlier showing that compelling evidence for the misinterpretation of garden-path sentences is found when participants engage in free recall allows us to set aside this concern. It is worth noting as well that the Tabor et al. findings provide support for the two-stage architecture embodied in the so-called Sausage Machine (Frazier & Fodor 1978), which assumes that a preliminary phrase packager (PPP) creates local parses and interpretations, which then must be put together by the sentence structure supervisor (SSS). The Tabor et al. results, then, can be viewed as indicating that reconciling the output of the PPP and the SSS is not a trivial task for the parser, perhaps because the parser tries to construct good enough, local parses rather than detailed and complete global structures containing long-distance dependencies.

In studies currently being conducted in our laboratory (F. Ferreira & J. D. Fodor, forthcoming) we have also found results that may indicate effects of mere local coherence. Participants were asked to read and then judge the grammaticality of globally ungrammatical sentences such as *While Katie fixed the car hit a fire hydrant*, and these were compared to the processing of similar grammatical items such as *When the gardener bathes his poodle joins him*. Eye movements were monitored. The ungrammatical sentences were incorrectly judged to be grammatical on about one-third of trials, which was much higher than for ordinary ungrammatical sentences. In addition, participants focused their reprocessing efforts on the verb *fixed* and the NP following (*the car*), suggesting that the system realizes that it is the problematic local sequence. Thus, the local acceptability of *Katie fixed the car* seems to interfere with processing of the sentence's global ungrammaticality. Again, these sorts of findings indicate that the comprehension system

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

works by cobbling together local analyses, and therefore has difficulty when local analyses cannot be reconciled with or integrated into a global structure.

The use of heuristics during parsing has received new support from the recording of ERPs during reading (van Herten et al. 2005). Participants encountered sentences that were semantically implausible – for example, one expressed the idea that a fox shot a poacher. Normally, the N400 response is sensitive to degree of semantic fit (Kutas & Hillyard 1980a,b,c, 1984), but in this study, the implausibility triggered a P600 – an index of syntactic revision. The authors argue that the system computed the semantically more compelling meaning of the sentence – that the poacher shot the fox – and that is why no N400 was observed. When later parts of the sentence conflicted with the syntactic skeleton for that meaning, a P600 occurred, indicating the system's attempt to reconsider the content of the sentence to make it fit the meaning derived from the semantic heuristic. Additionally, just as Tabor et al. (2004) found, locally plausible sequences interfered with overall global processing. The authors conclude that the comprehension system indeed employs a plausibility heuristic, and the conflict between the output of the heuristic and the output of syntactic algorithms gives rise to the P600.[2] Similar findings have been reported by others in the ERP community (Osterhout & Holcomb 1992; Osterhout & Mobley 1995; van Berkum et al. 2005; Kim & Osterhout 2005).

A great deal of work, then, supports the claim that people tend to misinterpret implausible sentences so they convey sensible propositions. An important reason why this might be so is that speakers sometimes do not say what they mean – that is, they make mistakes during production. For example, people sometimes make word exchange errors and say things like *I left the briefcase in my cigar* (Garrett 1976). If the comprehension system were not somewhat forgiving, it would not appreciate that a speech error occurred and that the right meaning is that the cigar was left in the briefcase, not the other way around. And, of course, the trigger for postulating a speech error is that the meaning of the utterance as spoken does not make sense. But, of course, people sometimes do need to express a strange and counterintuitive idea, so it is good that the parser does not always normalize sentence content. Nonetheless, there is growing evidence that the system computes a sensible meaning and then tries to reconcile it with the sentence's form. If the two conflict, then the parser appears to try to restructure the sentence, and then either goes with the altered structure or adjusts the meaning to conform to the structure and to allow expression of the unexpected idea.

Thus, people make errors when they talk, and listeners know speakers are fallible (after all, they are speakers too and they know they make mistakes). We might expect, then, that the comprehension system has tools for handling these deviations from ideal delivery.[3] One mechanism just discussed might be to normalize sentences that are implausible. Another set of tools might be devoted to handling utterances containing disfluencies. A

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

disfluency is any item in an utterance that causes it to deviate from an ideal delivery. Disfluencies are common – they are estimated to occur about six times per 100 words (Fox Tree 1995). Examples include fillers such as *um* and *er*, repeated words (*hand me the the pliers*), long, prosodically unmotivated pauses, and corrections (*hand me the the pliers um the tweezers*). Psycholinguists have generally ignored the question of how people process disfluencies in real time, but as was pointed out in the 2002 article, this question falls naturally out of the GE approach. Moreover, if we are to develop a complete theory of language understanding that takes into account how the system must have evolved given the nature of the input it has adapted to receive, then it becomes clear that an absolutely pressing question is how the parser handles disfluencies.

Ferreira and colleagues have been actively engaged in exploring this question. Bailey and Ferreira (2003) reported that fillers affect the normal operation of the parser. Recall the garden-path sentence we provided earlier, *When the gardener bathes his poodle joins him*. This garden-path sentence is even harder to process when the disambiguator *joins* and the head of the ambiguous phrase *poodle* are separated lexically, as in *When the gardener bathes his poodle with the soft fur joins him* (Ferreira & Henderson 1991). This effect occurs because the ambiguous phrase is assigned a specific syntactic and semantic role at its head, and the further the head from the disambiguator, the longer the parser has been committed to the wrong analysis of the ambiguous phrase. Bailey and Ferreira found the same effect with disfluencies – that is, they found precisely the same pattern of results when fillers occurred between *poodle* and *joins* as was observed with lexical items in the same position. It appears, then, that one effect that disfluencies have on comprehension is that fillers allow the parser to strengthen the syntactic and semantic analysis it has built up to that point. In addition, Bailey and Ferreira found some evidence for a cuing function: When people heard a filler before a phrase, which could be either simple or complex, they tended to assume the phrase would end up being complex. This tendency may be attributable to the parser's knowledge of co-occurrence frequencies between fillers and syntactic types. People are more likely to need to pause and say *um* or *er* when they are about to say something difficult (Goldman-Eisler 1968), so fillers tend to precede long and complex constituents.

Other work has examined how people process corrections in speech (Bailey 2004; Ferreira et al. 2004; Lau & Ferreira 2005), as in *Turn left I mean turn right at the light*. Some terminology will be helpful at this point: the part of the utterance spoken in error (*left* in the example) is typically called the reparandum, and the part produced to correct the reparandum is the repair (*right* in the example). The important result is that the comprehension system does not completely erase the reparandum from its representation of the utterance. Ferreira et al. (2004) asked participants to listen to sentences such as *I want you to put uh drop the ball* and then judge their grammaticality. Ferreira et al. found that people judged sentences like this one to be

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

grammatical less often than a version without the disfluency. They argued that the reparandum *put* lingers in the representation, which leads the parser to expect the verb *drop* to occur with a prepositional phrase (e.g., *drop the ball in the cup*) because *put* must have a prepositional phrase argument. When no such phrase is found in the input, the sentence is judged to be bad. The experiments also showed that *I want you to drop uh put the ball* was judged grammatical **more** often than a version without a disfluency, which is explained by the complementary mechanism: the repair *put* benefits from the lingering representation of the reparandum *drop*, and *drop* does not require a prepositional phrase argument. This latter finding is important because it rules out the idea that a sentence might be judged to be less acceptable merely because it contains a disfluency. This idea is contradicted by the finding that *I want you to drop uh put the ball* is judged to be better than a version with no disfluency (again, because the reparandum biases the interpretation of *put*). Similarly, Lau and Ferreira (2005) observed that the reduced relative structure could be made easier if it included the right sort of correction. They had participants judge the grammaticality of spoken sentences such as *The girl chosen uh picked for the play celebrated with her family and friends*. They found that the sentence was easier when it included a reparandum in the unambiguous past participle form, indicating again that the reparandum plays a role in overall comprehension.

One concern, of course, is that these studies on the comprehension of utterances containing corrections used off-line grammaticality judgments, so they do not provide information about how the reparandum is processed online, as it is encountered. Bailey (2004) remedied this flaw. Bailey used the so-called 'visual world paradigm' to explore how corrections are processed incrementally. Participants heard sentences such as *Hop the shark er the frog two spaces to the left*, and the listeners' task was to perform the appropriate action. Their eye movements were monitored with a free-viewing eye-tracking device. Because people tend to look at an object before moving it, patterns of fixations and saccades provide information about the interpretation that is created as the utterance unfolds. Bailey found that looks to the shark persisted long past the repair part of the correction (compared to a variety of control conditions), demonstrating that the reparandum indeed lingers in the representation. Very rarely did a participant mistakenly pick up the object denoted by the reparandum (i.e., the shark in this example), but their eye movements betrayed that at least briefly they entertained the meaning that the shark was the thing to be hopped. This occurred even though participants were told before the experiment that the animals' actions were constrained, so that (for example) only frogs could hop, only sharks could swim, and only soldiers could march. Thus, we see that the comprehension system attempts to suppress the reparandum, indicating that it indeed has mechanisms for handling disfluencies. At the same time, the process is just 'good enough' but not flawless, as the reparandum is not inhibited completely. Ferreira et al. (2004) describe a mechanism they

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

term 'overlay' to explain this behaviour on the part of the comprehension system.

New Directions

Good enough language comprehension is a proposal about how little information the comprehension system typically consults when it tries to derive an interpretation, and how shallow and even wrong the resulting interpretation can therefore sometimes be. But it is important to note that most of the time, the system will either compute the correct meaning or any discrepancies from what was intended will not be significant from the perspective of successful communication. To put this issue into perspective, it is important to appreciate that when we say a sentence has been 'misunderstood', we mean that the representation is incorrect in some relevant way; we do not mean that it is completely unrelated to its actual content. For example, consider once again the sentence *While Mary bathed the baby played in the crib*. It is true that most participants think this sentence means that Mary bathed the baby, and that that interpretation is incorrect.[4] Nonetheless, it is important to give people credit for what they do understand. They appreciate that there is a bathing action involving Mary as the agent, and a playing action involving a baby in a crib. They might even understand that Mary bathed herself (Christianson et al. 2006) at the same time that they think Mary bathed the baby. Thus, they have more of the sentence right than wrong. Moreover, the misinterpretation might not harm communication, as people often do bathe babies.

What these experiments on GE processing have done, then, is to set up semi-artificial situations that allow language scientists to uncover the hidden workings of the system. The result is that we have been able to see the shallow and even incorrect representations that sometimes get created. In this way, this research programme is much like the one in decision-making that attempts to uncover the heuristics and biases that people use when they reason (Tversky & Kahneman 1974). The idea is that people **seem** to reason effectively or seem to understand utterances completely and correctly, but to a large extent this appearance is based on the skillful application of some fairly superficial heuristics.

We would argue that the next step in this research programme on GE comprehension is similar to what has happened recently in the literature on decision-making: rather than continue to focus on demonstrating shallow and error-prone processing, we might want to explore why the application of heuristics is in fact adaptive, as the ABC community has advocated (Gigerenzer et al. 1999). Gigerenzer and colleagues have argued that a cognitive system that considers all potentially relevant sources of information to arrive at the 'best' decision is not biologically realistic, as humans are constrained by limits of time and resources. In the terminology of the ABC group, only 'demons' could reason in this way. Gigerenzer et al. have

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

demonstrated that a system that is resource limited and that must make decisions quickly actually does better if it relies on a small set of 'fast and frugal' heuristics rather than attempting to execute algorithms that consult every potentially relevant piece of information.

Given this perspective, it is interesting to return to models of parsing like the garden-path model, in which it is assumed that a simple heuristic governs all parsing decisions: minimize all processing steps (see Fodor & Inoue's 'minimal everything' principle, 1998). That is, the minimal attachment principle states that the parser should build the simplest structure it can, and revise it only if necessary. The late closure heuristic leads the parser to attach material to what is currently in its domain of processing, rather than trying to retrieve an old constituent or setting up a new one. The minimal chain principle assumes that the parser tries to make the distance between a moved constituent and its trace or gap as short as possible. The revision-as-last-resort principle leads the parser to change what it has done only if there is something wrong with the representation that it has built, not just because the representation could be better. In other words, if it isn't broken, don't fix it. And, finally, the minimal revisions principle (Fodor & Inoue 2000) says that if the parser does engage in repair, it should make the fewest changes possible to fix the problem, and should not attempt to create a globally ideal representation. Thus, the garden-path model is a minimal effort approach. The architecture allows the comprehension system to quickly build a representation that can support some type of interpretation. If later the system encounters a road block, then it will engage in revision, but it will do as little as possible to get back on track.

This type of architecture seems well suited to the demands of real-world communication. One important point that has not been appreciated is that understanding and revising sentences takes time, which might not be available during dialogue. For some difficult garden-path sentences, people will linger for longer than a second on the last word or two if they control the presentation of the input. But in real conversations, we do not control how utterances are presented to us. If we were to spend a second or two engaging in re-analysis, end-of-sentence wrap-up, and all the other processes that take measurable time, we would fall further and further behind the input, and communication would indeed be affected. It seems more likely that new material preempts processing of what has already been received, particularly when a 'good enough' representation has already been built.

Another intriguing lesson to learn from the comparison to the literature on decision-making is that models which postulate that all interpretations are activated in parallel and that all sources of information interact freely to determine which interpretation eventually wins (Trueswell et al. 1993; MacDonald et al. 1994; Stevenson, 1994; Spivey & Tanenhaus 1998) seem to assume unbounded rationality (Gigerenzer & Selten 2001). They require a great deal of bandwidth to allow more than one representation to be maintained and assessed, and to allow everything from co-occurrence

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

frequency to prosody to real-world plausibility to speaker intentions to influence the strength of activation for each representation. If the demon approach is not realistic for decision-making, it is at least conceivable that it is not realistic for language comprehension either. On this view, then, GE representations arise because the system makes use of a set of heuristics that allow it to do the least amount of work necessary to arrive at a meaning for a sentence. If this suggestion is correct, then the direction for further research should be to determine whether factors such as time pressure and resource constraints do indeed influence the extent to which the system operates in a GE way. Our hope is that this article encourages researchers in psycholinguistics to begin moving in this research direction.

Short Biography

Fernanda Ferreira obtained her PhD from the University of Massachusetts at Amherst in 1988. She is now Professor and Chair of Language and Cognition at the University of Edinburgh. Her research areas are language comprehension and production, focusing especially on the roles of syntax, prosody, and disfluencies, and the integration of visual and linguistic information online. She is co-author with John M. Henderson of *The Integration of Language, Vision, and Action* (2004, Psychology Press), and co-author with Janet D. Fodor of *Reanalysis in Sentence Processing* (1999, Kluwer).

Nikole Patson's area of research is psycholinguistics. Current research focuses on plural noun representation and lingering misinterpretations. She holds a BS in Psychology from the University of Michigan-Flint.

Endnotes

⋆ Correspondence address: Fernanda Ferreira, Psychology, 7 George Square, University of Edinburgh, Edinburgh EH8 9JZ, UK.

[1] The study also demonstrates that the tendency to misinterpret depends both on the comprehender's working memory capacity and the delay between presentation of the sentence and recall, but these are topics beyond the scope of this article.

[2] An important issue that has not been addressed directly in the literature on 'good enough' language processing is whether mere implausibility is processed differently than semantic anomaly, as has been found in recent studies employing the eye movement monitoring methodology (e.g., Rayner et al. 2004).

[3] One might object that most studies showing normalization have been conducted with written materials, and errors are far less common in writing than in speech. There are several responses to this concern. One is that some studies have been conducted with spoken language – for example, Ferreira (2003). Second, although written materials are less likely to contain errors because people have the opportunity to revise privately, many errors still make it through, as anyone who has corrected student papers knows. Finally, although the listening and reading systems are different, many processes overlap, and those that are the same might result in reading behavior similar to that observed in listening contexts.

[4] Note that we do not mean that it makes no sense to think *Mary might have bathed the baby*; what we mean is that that is not what the sentence says. The important point is that the interpretation is more likely to include the proposition that *Mary bathed the baby* if the syntactic structure of the

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

sentence leads to a garden path. Thus, if the two clauses are presented in reverse order or separated by a comma, the misinterpretation almost never occurs.

Works Cited

Bailey, K. G. B. 2004. Disfluent speech and the visual world: An application of the visual world paradigm to the study of spoken language comprehension. Unpublished Doctoral Dissertation, Michigan State University, East Lansing, MI.

Bailey, K. G. B., and F. Ferreira. 2003. Disfluencies influence syntactic parsing. Journal of Memory and Language 49.183–200.

Barton, S. B., & A. J. Sanford. 1993. A case study of anomaly detection: shallow semantic processing and cohesion establishment. Memory and Cognition 21.477–87.

Bever, T. G. 1970. The cognitive basis for linguistic structures. Cognition and Language Development, ed. by R. Hayes, 279–362. New York: Wiley & Sons, Inc.

Christianson, K., A. Hollingworth, J. Halliwell, and F. Ferreira. 2001. Thematic roles assigned along the garden path linger. Cognitive Psychology 42. 368–407.

Christianson, K., C. Williams, R. Zacks, and F. Ferreira. 2006. Younger and older adults' good enough interpretations of garden-path sentences. Discourse Processes 42.205–38.

Erickson, T. A., and M. E. Matteson. 1981. From words to meanings: a semantic illusion. Journal of Verbal Learning and Verbal Behavior 20.540–552.

Ferreira, F. 2003. The misinterpretation of noncanonical sentences. Cognitive Psychology 47.164–203.

Ferreira, F., and J. M. Henderson. 1991. Recovery from misanalyses of garden-path sentences. Journal of Memory and Language 31.725–45.

Ferreira, F., K. Christianson, and A. Hollingworth. 2001. Misinterpretations of garden-path sentences: implications for models of reanalysis. Journal of Psycholinguistic Research 30.3–20.

Ferreira, F., V. Ferraro, and K. G. D. Bailey. 2002. Good-enough representations in language comprehension. Current Directions in Psychological Science 11.11–15.

Ferreira, F., E. F. Lau, and K. G. D. Bailey. 2004. Disfluencies, parsing, and tree-adjoining grammars. Cognitive Science 28.721–49.

Fodor, J. D., and A. Inoue. 1998. Attach anyway. Reanalysis in sentence processing, ed. by J. D. Fodor and F. Ferreira, 101–141. Dordrecht, The Netherlands: Kluwer.

Fodor, J. D., and A. Inoue. 2000. Garden path reanalysis: attach (anyway) and revision as last resort. Cross-Linguistic Perspectives in Language Processing, ed. by M. DiVincenzi and V. Lombardo, 21–61. Dordrecht, The Netherlands: Kluwer.

Fodor, J. D., and F. Ferreira. (forthcoming). Processing ungrammatical sentences: evidence from eye movements and speeded grammaticality judgments.

Fox Tree, J. E. 1995. The effects of false starts and repetitions on the processing of subsequent words in spontaneous speech. Journal of Memory and Language 34.709–38.

Frazier, L. 1979. On comprehending sentences: Syntactic parsing strategies. Doctoral Dissertation, University of Connecticut, Storrs, CT.

Frazier, L., and J. D. Fodor. 1978. The sausage machine: a new two-stage parsing model. Cognition 6.291–325.

Frazier, L., and K. Rayner. 1982. Making and correcting errors during sentence comprehension: eye movements in the analysis of structurally ambiguous sentences. Cognitive Psychology 14.178–210.

Frazier, L., J. M. Pacht, and K. Rayner. 1999. Taking on semantic commitments, II: collective versus distributive readings. Cognition 70.87–104.

Garrett, M. F. 1976. Syntactic processes in sentence production. New approaches to language mechanisms, R. J. Wales and E. Walker. Amsterdam: North-Holland.

Gigerenzer, G., and R. Selten. (Eds.) 2001. Bounded rationality: The adaptive toolbox. Cambridge, MA: MIT Press.

Gigerenzer, G., Todd, P. M., and the ABC Group. 1999. Simple heuristics that make us smart. New York: Oxford University Press.

Goldman-Eisler, F. 1968. Psycholinguistics: Experiments in spontaneous speech. London: Academic Press.

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd

Kim, A., and L. Osterhout. 2005. The independence of combinatory semantic processing: evidence from event-related potentials. Journal of Memory and Language 52.205–25.
Kutas, M., and S. A. Hillyard. 1980a. Event-related brain potentials to semantically inappropriate and surprisingly large words. Biological Psychology 11.99–116.
Kutas, M., and S. A. Hillyard. 1980b. Reading between the lines: event-related brain potentials during natural sentence processing. Brain and Language 11.354–73.
Kutas, M., and S. A. Hillyard. 1980c. Reading senseless sentences: brain potentials reflect semantic incongruity. Science 207.203–5.
Kutas, M., and S. A. Hillyard. 1984. Brain potentials during reading reflect word expectancy and semantic association. Nature 307.161–3.
Lau, E., and F. Ferreira. 2005. Lingering effects of disfluent material on comprehension of garden path sentences. Language and Cognitive Processes 20.633–66.
MacDonald, M. C., N. J. Pearlmutter, and M. S. Seidenberg. 1994. The lexical nature of ambiguity resolution. Psychological Review 101.676–703.
Osterhout, L., and P. J. Holcomb. 1992. Event-related brain potentials elicited by syntactic anomaly. Journal of Memory and Language, 31.785–806.
Osterhout, L., and L. A. Mobley. 1995. Event-related brain potentials elicited by failure to agree. Journal of Memory and Language 34.739–73.
Patson, N. D., E. Swensen, N. Moon, and F. Ferreira. 2006. Individual differences in syntactic reanalysis. Poster presented at the Annual CUNY Conference on Human Sentence Processing, March 2006, New York, NY.
Pickering, M. J., and S. Frisson. 2001. Processing ambiguous verbs: evidence from eye-movements. Journal of Experimental Psychology: Learning, Memory *, and* Cognition 27.556–73.
Rayner, K., T. Warren, B. Juhasz, and S. P. Liversedge. 2004. The effect of plausibility on eye movements in reading. Journal of Experimental Psychology: Learning, Memory, & Cognition 30.1290–301.
Sanford, A. J., and P. Sturt. 2002. Depth of processing in language comprehension: not noticing the evidence. Trends in Cognitive Sciences 6.382–6.
Spivey, M., and M. Tanenhaus. 1998. Syntactic ambiguity resolution in discourse: modeling the effects of referential context and lexical frequency. Journal of Experimental Psychology: Learning, Memory, and Cognition 24.1521–43.
Stevenson, S. 1994. Competition and recency in a hybrid network model of syntactic disambiguation. Journal of Psycholinguistic Research 23.295–322.
Tabor, W., B. Galantucci, and D. Richardson. 2004. Effects of merely local syntactic coherence on sentence processing. Journal of Memory and Language 50.355–70.
Townsend, D., and T. G. Bever. 2001. Sentence comprehension: The integration of habits and rules. Cambridge, MA: MIT Press.
Trueswell, J. C., M. K. Tanenhaus, and C. Kello. 1993. Verb-specific constraints in sentence processing: separating effects of lexical preference from garden-paths. Journal of Experimental Psychology: Learning, Memory, and Cognition 19.528–53.
Tversky, A., and D. Kahneman. 1974. Judgment under nncertainty: heuristics and biases. Science 185.1124–31.
Van Berkum, J. J. A., C. M. Brown, P. Zwitserlood, V. Kooijman, and P. Hagoort. 2005. Anticipating upcoming words in discourse: Evidence from ERPs and reading times. Journal of Experimental Psychology: Learning, Memory, & Cognition 31.443–67.
Van Herten, M., H. H. J. Kolk, and D. J. Chwilla. 2005. An ERP study of P600 effects elicited by semantic anomalies. Cognitive Brain Research 22.241–55.

© 2007 The Authors
Journal Compilation © 2007 Blackwell Publishing Ltd