Received: 7 September 2020 | Revised: 4 June 2021 | Accepted: 5 July 2021

DOI: 10.1111/lnc3.12434

ARTICLE

WILEY

Why do readers fail to notice word transpositions, omissions, and repetitions? A review of recent evidence and theory

Kuan-Jung Huang | **Adrian Staub**

Department of Psychological and Brain Sciences, University of Massachusetts Amherst, Amherst, Massachusetts, USA

Correspondence
Adrian Staub, Department of Psychological and Brain Sciences, University of Massachusetts Amherst, 430 Tobin Hall, Amherst, MA, 01003, USA.
Email: astaub@umass.edu

Funding information
US National Science Foundation, Division of Behavioral and Cognitive Sciences, Grant/Award Number: BCS 1732008

Abstract

Most readers have had the experience of initially failing to notice an omission or repetition of a function word, or a transposition of two adjacent words. In the present article, we review recent research investigating this phenomenon. We emphasize that failure to notice such errors is of substantial theoretical interest, given what we have learned about how systematically and incrementally readers inspect and process text. We endorse the idea that a process of rational inference may play a critical role, while we cast doubt on the idea that failure to notice errors arises from parallel processing of multiple words. We review a number of recent studies from our own laboratory that have investigated the relationship between eye movements during reading and noticing, or failing to notice, an error. While the conclusions from these studies are broadly consistent with a rational inference account, we find that when readers fail to notice an error, their eye movements generally show no indication that the error was registered at all. On its surface, this finding may be viewed as inconsistent with the idea that the rational inference process that enables readers to overlook errors is genuinely post-perceptual. We suggest a mechanism by which eye movement control models could account for this finding.

1 | INTRODUCTION

We begin by asking you to read the opening paragraph of Ernest Hemingway's classic novel The Old Man and the Sea:

> He was an old man who fished alone in a skiff in the Gulf Stream and he had gone eighty-four days now without taking a fish. In the first forty days a boy had been with him. But after forty days without a fish the the boy's parents had told him that the old man was now definitely and finally *salao*, which is the worst form of unlucky, and the boy had gone at their orders in another boat which caught three good fish the first week. It made the boy sad see the old man come in each day with his skiff empty and he always went down to him help carry either the coiled lines or the gaff and harpoon and the sail that was furled around the mast. The sail was patched with flour sacks and, furled, it looked like the flag of permanent defeat. (Hemingway, 1952)

With apologies to Hemingway, we have inserted three errors in the text. In the third sentence, the word *the* is repeated. In the fourth sentence, the word *to* is missing between the words *sad* and *see*, and the words *him* and *help* are transposed. If you are like many skilled readers, you failed to notice one or more of these errors.

If you have ever done any proofreading, you are probably already aware that such errors are easily overlooked. Until quite recently, however, this phenomenon has not elicited much scientific investigation, but has instead been the province of informal demonstrations and internet memes. For an amusing example, search on the internet for the phrase *What I if told you you read that wrong*. Another example is included in many introductory psychology textbooks (e.g., Eysenck & Keane, 2005), where the phrase *Paris in the the spring* is printed in a triangle, with a line break between the two instances of 'the'. The authors of these textbooks predict—correctly, it seems—that most readers will initially fail to notice the repetition, but they provide only extremely informal explanation in terms of 'top-down' processing.

This situation has now changed, as researchers have turned their attention to the question of how frequently readers fail to notice word omissions, transpositions, or repetitions, and why they do. In this article, we review this research. We discuss the theoretical explanations that are on offer, as well as empirical and theoretical challenges for these accounts. We proceed as follows. In the next section, we argue that readers' failure to notice these errors is indeed surprising and theoretically interesting, given what we know about how incrementally and systematically fluent readers extract information from the words on the printed page. In the following sections we discuss several lines of recent research, including our own, that elucidate different aspects of this phenomenon.[1]

2 | SHOULD WE BE SURPRISED BY READERS' FAILURE TO NOTICE ERRORS?

On one view of the reading process, encapsulated in the title of Kenneth Goodman's (1967) famous article Reading: A Psycholinguistic Guessing Game, it is not unexpected that readers should sometimes fail to notice errors. Goodman disputed the notion that skilled reading results 'from precise perception and identification of all elements', arguing instead that it involves 'selecting the fewest, most productive cues necessary to produce guesses which are right the first time' (p. 127). He asserted that the development of reading skill 'involves not greater precision, but more accurate first guesses based on better sampling techniques....As the child develops reading skill and speed, he uses increasingly fewer graphic cues' (p. 132). Anticipating the present topic, Goodman noted that if readers do systematically process each word, the situation of 'a highly proficient adult reader reading and rereading a paper he's written and always missing the same misprints' (p. 131) is quite mysterious. But if, as Goodman maintained, the reader makes guesses about the content of the text that are only minimally constrained by the visual input, it is unsurprising that she will fail to notice errors such as word omissions or repetitions.

However, the intervening half century of reading research has shown Goodman's conception of the reading process to be mostly, if not entirely, incorrect. The large body of research in which readers' eyes are tracked on a millisecond-by-millisecond basis (see Rayner, 1998 for review) has revealed that skilled readers directly fixate most of the words on the page (Brysbaert & Vitu, 1998); that the time they spend on a word is reliably influenced by variables such as the word's frequency of occurrence (e.g., Rayner & Duffy, 1986; Staub et al., 2010); that when the eyes do skip a word, the word is processed to a fairly high level prior to the skip (e.g., Choi & Gordon, 2014); and that when a word is syntactically or semantically anomalous, readers tend to slow down almost immediately, or make regressive eye movements to an earlier point in the text (e.g., Frazier & Rayner, 1982; Rayner et al., 2004). A reader's expectations based on linguistic and real-world knowledge do play a role in how they inspect and process the text, but this role is highly constrained. For example, though the probability that readers' eyes skip a word is influenced by the word's predictability (as measured by the proportion of participants who provide that word in a sentence-completion or cloze task), they do so only when the predictable word is visible in the parafovea, to the right of their current eye fixation; when some other word or letter string appears in place of the predictable word, readers are not likely to skip it (Balota et al., 1985). While readers may 'guess' the identity of an upcoming word, they do not skip the word on this basis alone—the bottom-up input in parafoveal vision has to confirm the reader's guess. Thus, our modern picture of the reading process implies, contra Goodman's view, that there is indeed a puzzle about why skilled readers fail to notice errors such as word omissions and repetitions.

Another view would simply attribute failures to notice errors to occasional episodes of 'zoning out' while reading, during which the reader allows her mind to wander from the primary task of text comprehension (Franklin et al., 2011; Reichle et al., 2010). Schad et al. (2012) tested this idea by monitoring eye movements while subjects read stories in which errors of various kinds had been inserted. Schad et al. observed that when a reader failed to detect an error, eye movements on the preceding words showed reduced effects of linguistic variables such as word frequency, suggesting that readers are in a different cognitive state on occasions when they detect an error and on occasions when they do not. However, Schad et al. investigated the detection of only highly salient errors, such as transposition of two words from very far apart in a sentence (e.g., *On all classmates, he congratulates his birthdays and the teachers*). Below we discuss a number of recent studies showing that readers tend to miss omissions or repetitions of function words, and transpositions of adjacent words, even when sentences are presented one at a time, and when the task encourages the reader to focus specifically on error detection. While mind wandering or 'zoning out' may play some role in failure to detect errors, it appears that some kinds of errors are missed even when readers are paying close attention.

It may also be pointed out that there are cases in which language comprehenders seem to either miss or hallucinate a word even in non-reading situations, so perhaps it is not necessary, in our attempt to understand the source of these failures, to focus on the role of reading as the input modality. However, these demonstrations clearly involve mechanisms that are distinct from those involved in normal reading. One example is repetition blindness (RB), first observed by Kanwisher (1987; Kanwisher & Potter, 1990). Kanwisher found that when a list of words is presented in rapid serial visual presentation (RSVP) format at less than 200 ms/word, subjects very frequently fail to recall a second instance of a word, reporting only the first instance. Remarkably, this was also the case when the stimuli were sentences, rather than word lists, and when the second instance was fully grammatical, for example, *When she spilled the ink there was ink all over*. Indeed, in Kanwisher's (1987) original experiment with sentences, subjects failed to recall the second instance of a word almost 80% of the time at a presentation rate of 117 ms/word, despite the fact that this resulted in report of an ungrammatical sentence, for example, *When she spilled the ink there was all over*. The specific circumstances under which RB arises have since been explored extensively, as have theoretical explanations in terms of perception or memory (e.g., Bavelier & Potter, 1992; Harris & Morris, 2000). For present purposes, however, what is most relevant is that failure to notice repeated words in reading is not easily assimilated to RB. First, RB for words does not arise at rates of presentation that more closely approximate normal reading; Kanwisher (1987) found no RB at all when stimuli were presented for 250 ms/word, which happens to be very close to the average reading rate (Brysbaert, 2019). In addition, RB is pronounced even when the two instances of a word are not adjacent, as was the case in the example sentence above, and RB can result in recalling a fully grammatical sentence as if it were ungrammatical. We have no evidence of analogous effects in normal reading.

Another very interesting phenomenon involving both missing and hallucinating words in auditory language comprehension was recently investigated by Dilley and Pitt (2010; Morrill et al., 2014). In Dilley and Pitt (2010), subjects heard sentences such as *Jill got quite mad when she heard there are birds* or *Jill got quite mad when she heard their birds*, which are both grammatical, but which differ in whether a critical function word (are) is present. Dilley and Pitt found that when the surrounding context was spoken at a slow rate relative to the stretch of speech around the critical word, listeners often failed to report the function word in the first sentence; that is, they took it to be the second sentence. However, when the surrounding context was spoken at a fast rate, listeners often reported a function word that was actually absent; that is, they took the second sentence to be the first. Dilley and Pitt hypothesized that the rate of the surrounding speech affects an early processing stage at which the acoustic input is segmented into words. However, in reading of an alphabetic language like English, word boundaries are unambiguously marked by spaces, which means that it is unlikely that an analogue of this effect is responsible for readers' failure to notice a word omission, or tendency to interpret two instances of a word as a single instance.

3 | ARE ERRORS PERCEIVED, BUT THEN CORRECTED BY INFERENCE?

The first perspective on this problem that we will consider is associated with Gibson and colleagues (Gibson et al., 2013; Ryskin et al., 2018; see also Levy, 2008), who have developed a 'rational inference' framework, often referred to as the noisy channel theory, for understanding why language comprehenders—either listeners or readers—sometimes misrepresent the input, and do so in systematic ways.[2] This account emphasizes that sentence comprehension involves recovering an intended message that may have been corrupted by various sources of noise during transmission, including by errors on the part of the producer or the perceiver. Optimal Bayesian inference by the comprehender will estimate the probability of an intended message, given the perceived message (P(Si/Sp); Si is the intended message, and Sp is the perceived message), as a function of the prior probability of the intended message (P(Si)), and how likely it is that given Si, the message would appear as perceived (P(Sp/Si)). The key predictions of this account are that comprehenders will misrepresent the input when (a) there is a potential error-generating mechanism that could have deformed an intended message so as to give rise to the perceived message (i.e., P(Sp/Si) is high); and (b) the intended message that is an alternative to the perceived input is quite likely (i.e., P(Si) is also high).

Illustrating the importance of P(Sp/Si), Gibson et al. (2013) found that readers are likely to reach a plausible interpretation by 'inserting' the word 'to' in a sentence such as *The mother gave the candle the daughter*, but are relatively unlikely to 'delete' the word 'to' in a sentence such as *The mother gave the daughter to the candle*. Gibson et al. suggest that the comprehender knows—at some level—that it is relatively likely that a producer would accidentally omit an instance of *to*, or that the comprehender herself would fail to notice an instance of *to* that was actually present; but the comprehender also knows that it is relatively unlikely than that the producer would accidentally insert an extra *to*, or that the comprehender herself would hallucinate an instance of *to* that was absent. Thus, the comprehender may infer the presence of *to* when reading a sentence with a missing *to*, but is relatively unlikely to ignore the presence of *to* in a sentence that should not include this word. Similarly, Poppels and Levy (2016) found that participants are fairly likely to infer a nonliteral, plausible meaning of a sentence like *The ball fell from the floor to the table*, while they are relatively unlikely to do so with *The ball kicked the boy*. Poppels and Levy note that speakers make exchange errors with prepositions relatively frequently, and as a result, the perceived input *from the floor to the table* is not particularly unlikely when the intended message is *to the floor from the table*.

On the other hand, the importance of P(Si)—the probability of an intended message that is not the perceived message—was also demonstrated by Gibson et al. (2013), who showed that when many filler sentences in an experiment are implausible, participants are less likely to 'correct' implausible target sentences to a nonliteral interpretation than when few filler sentences are implausible. It appears that participants take into consideration the probability of encountering a plausible or implausible sentence locally within an experiment. The higher the base rate of implausible sentences, the less likely it is that the reader will consider an implausible target sentence to be the result of a communication mistake. An even more striking demonstration of the role of P(Si) is from a recent study by Keshev and Meltzer-Asscher (2021). In this study, readers of Hebrew seem to 'hallucinate' subject-verb agreement errors in a sentence that actually has correct agreement, when the correct agreement would require the reader to adopt a syntactic analysis on which the sentence has a very unusual structure.

The interpretation of these studies is sometimes complicated by the fact that the experimental tasks may have implicitly encouraged inference. For example, in Gibson et al. (2013) subjects answered questions such as *Did the daughter receive something?* after the sentence *The mother gave the candle the daughter*. It is uncertain whether a subject who answers 'yes' to this question has engaged in a rapid, automatic form of Bayesian inference, resulting in her failing to notice the absence of *to*, or is simply responding to perceived task demands; see Ryskin et al. (2018) for a version of the task that even more explicitly encourages deliberate inference about what the speaker might have meant. However, there is also evidence for inference occurring during online sentence processing (Keshev & Meltzer-Asscher, 2021; Levy et al., 2009; Levy, 2011), without any relevant task demands.

In sum, the noisy channel theory, as described by Gibson and colleagues, proposes that readers (and listeners) overlook certain kinds of errors not because they do not perceive them, but because they perceive them and then 'correct' them; a veridical representation of the input Sp is directly involved in noisy channel inference. As Gibson et al. (2013) state, 'The comprehender thus observes Sp and must decode it to its intended meaning Mp' (p. 8051). Or, as Ryskin et al. (2018) state, '[I]f comprehenders perceive an implausible sentence Sp (e.g., The oven cleaned the grandmother) which is "close" to a more plausible sentence (e.g., The grandmother cleaned the oven), they should infer that the producer actually uttered (or intended to utter) the plausible version.' (p. 142). However, the claim that readers do initially form a veridical representation of the input, including any errors it contains, may be regarded as quite puzzling, in light of the fact that readers often appear to be completely unaware of errors. For example, we assume that you may have been completely unaware of one or more of the errors in the Hemingway passage that began this article, on your first reading of this passage. If this is right, then the noisy channel account must allow that the inference process, even if post-perceptual, is so rapid and automatic that it prevents perceived errors from ever reaching awareness; all that reaches awareness is the 'corrected' message itself. We return to this point below.[3]

4 | WORD TRANSPOSITIONS AND PARALLEL LEXICAL PROCESSING

We now consider an entirely different theoretical perspective. In Mirault et al. (2018), participants were asked to make rapid grammaticality judgements to five-word French sentences; (1) provides examples of their three conditions in English translation.

1. a. The white cat was big.
   b. The white was cat big.
   c. The white was cat slowly.

In two experiments, subjects made errors on 12%–16% of trials with ungrammatical sentences like (1b), compared to about 3% with grammatical sentences like (1a) and ungrammatical sentences like (1c); they also responded more slowly to sentences like (1b). Critically, (1b) was created by transposing two words from a grammatical sentence, so a response of 'grammatical' would result from failing to notice the order of *was* and *cat*. A follow-up study by Snell and Grainger (2019b) used the same task with four-word sentences, and compared transposition of the two adjacent inner words versus the two outer words. The error rate and mean response time were significantly greater in the inner transposition condition.

Mirault et al. (2018) explained these results by deploying an idea that is not new (Engbert et al., 2005; Reilly & Radach, 2006), but which has been resuscitated by this research group in several recent papers (e.g., Snell et al., 2018; Snell & Grainger, 2017). The critical claim is that during normal reading, multiple words are processed at the same time. Thus, in (1b) *was* and *cat* are processed simultaneously, and on occasion *cat* may be recognized before *was*. To fully explain the failure to notice a transposition, Mirault et al. also propose that each word is not initially tagged to a definitive spatial location, and that knowledge about sentence structure can be brought to bear very rapidly, guiding the reader to represent *cat* as occurring to the left of *was*.

Parallel lexical processing has long been an attractive idea to some eye movement researchers (see Reichle, Liversedge et al., 2009, for discussion), but several lines of evidence have reduced its appeal. First, high-level parafoveal-on-foveal effects, which would seem to be predicted by parallel processing models, do not generally appear in eye movement experiments. If words are processed in parallel, properties of the word to the right of the fixated word, such as its frequency, should influence the duration of the current eye fixation. But though some corpus studies have revealed longer fixation times on a word when it is followed by a less frequent word (Kennedy & Pynte, 2005; Kliegl et al., 2006), most carefully controlled experimental studies have failed to find these effects; see Brothers et al. (2017) for a particularly comprehensive assessment. In light of these null findings, Snell and colleagues (Snell et al., 2017; Snell et al., 2018; Snell & Grainger, 2019a) have proposed that a spatiotopic sentence representation could enable readers to keep track of which information belongs to which word, enabling a parallel processing model to avoid predicting parafoveal-on-foveal effects. But a simpler explanation is that lexical processing is serial: A word is not lexically processed until the previous word has been recognized, as proposed by serial-attention-shift models such as E-Z Reader (Reichle et al., 1998).

A second challenge comes from recent work by White and colleagues (White et al., 2018; 2020; White et al., 2019), who address the fundamental question of whether readers can recognize two words at the same time, when the task is set up to require them to do so. Participants in their experiments perform a lexical decision or semantic categorization on two words presented simultaneously. In one condition, participants are cued to attend to one word only, while in another, they must divide attention between the two words. Critically, the stimuli are presented so briefly that the task is relatively difficult even when only one word is attended. Across many experiments, performance in the divided attention condition falls significantly below the prediction of models that assume parallel processing, either with or without capacity limitations, but aligns well with the prediction of an all-or-none serial model; it appears that on each trial participants are able to process one word or the other, but not both. A functional magnetic resonance imaging experiment using the same paradigm (White et al., 2019) provides converging evidence. The BOLD response in the anterior part of the left ventral occipitotemporal cortex (VWFA-2) showed sensitivity to the frequency of only one word, even when attention is divided.

A third challenge for parallel lexical processing as an explanation for readers' failure to notice errors is that it is unclear whether this account can explain failure to notice word repetitions or omissions. It is conceivable that parallel processing could result in two adjacent instances of the same word being encoded as a single instance, and indeed, there is some evidence that when the boundary paradigm (Rayner, 1975) is used to present a word in the parafovea, to the right of fixation, that is identical to the fixated word, processing of the fixated word is actually facilitated (Angele et al., 2013; Dare & Shillcock, 2013). This finding is consistent with the idea that orthographic information is being obtained from both words at the same time. But below, we discuss evidence that actual failure to notice repetitions depends very strongly the specific repeated word, with readers very frequently failing to notice a repetition of the word *the*, but very rarely failing to notice a repetition of a noun (Staub et al., 2019). Parallel lexical processing would not seem to offer an explanation for this sharp dissociation. In addition, it is not at all clear how failure to notice a word omission might arise due to parallel lexical processing. Thus, to the extent that we seek a unified explanation of these phenomena, parallel lexical processing would not appear to be a promising candidate.

In sum, while Mirault et al. (2018) provide a convincing demonstration that readers do sometimes overlook transposition errors, even in an explicit error detection task, their theoretical explanation in terms of parallel lexical processing faces substantial challenges, due to the apparent lack of high-level parafoveal-on-foveal effects, and due to the recent behavioural and neuroimaging work by White and colleagues. Moreover, the notion of parallel lexical processing appears to be most relevant to the failure to notice word transpositions; unlike noisy channel inference, it does not have obvious potential to provide a unified explanation of failure to notice various kinds of errors.[4]

5 | WHAT CAN EYE MOVEMENTS TELL US?

In our own laboratory, we have recently conducted a number of experiments using eye-tracking in studies of error detection during reading. Several previous eye movement experiments have investigated how readers' eyes respond when a phrase structure or agreement error is encountered in the course of reading (Ni et al., 1998; Pearlmutter et al., 1999). However, with the exception of the Schad et al. (2012) study mentioned above, none has explicitly probed whether subjects notice these errors, and none has attempted to identify eye movement behaviour associated with noticing an error or, perhaps more crucially, failing to notice it. In our work, we have attempted to answer four central questions: (1) How prevalent is failure to notice errors in a relatively natural reading situation, where error-detection is not subjects' central task? (2) How strong is any relationship between direct fixation on an error and the tendency to notice it? (3) Does an error disrupt a reader's eye movements even when it is not noticed? (4) Is there a relationship at the level of individual subjects between reading behaviour and the tendency to notice errors?

Staub et al. (2019) asked how often readers fail to notice a repetition or omission error. Each critical sentence had four variants: grammatical (G), repeated-the (RT), repeated-noun (RN), and omitted-the (OT), as shown in (2).

2. a. Amanda jumped off the swing and landed on her feet (G).
   b. Amanda jumped off the the swing and landed on her feet (RT).
   c. Amanda jumped off the swing swing and landed on her feet (RN).
   d. Amanda jumped off swing and landed on her feet (OT).

The 36 critical sentences were intermixed with 72 fillers. While fillers were followed by a comprehension question, critical sentences were followed by the question, 'Was there anything wrong with that sentence?' The goal of this intermixing of question types was to reduce the demand to treat the task as an explicit proofreading task. While participants were nearly perfect in accepting the grammatical sentences, and were about 90% accurate in rejecting the RN sentences, they noticed the error in the OT sentences only 67% of the time, and most strikingly, noticed the error in the RT sentences only 46% of the time.

The only other study that has explicitly investigated processing of fully-visible, repeated words (Healy & Zangara, 2017) also found that a repeated *the* is particularly problematic. The results of this study indicate, moreover, that the difficulty in noticing a repetition of the word *the* is not due entirely to its length or syntactic function. Healy and Zangara had subjects read aloud sentences with a repeated *the*, or sentences with the word *one* appearing in the same position (e.g., *Two cougars stalked the the/one one prairie dog.*). They found that subjects failed to correctly read aloud the two instances of *the* more than 30% of the time, while missing one of the instances of *one* less than 5% of the time.

The analysis of readers' eye movements in Staub et al. (2019) revealed several interesting patterns. First, when readers directly fixated both instances of *the* in the RT sentences, the detection rate did increase, but only to 66%; thus, word skipping is only one factor in failure to notice the repetition. Similarly, directly fixating the words before and after the position of the omitted *the* in the OT sentences (*off* and *swing* in 2d) resulted in only a very modest increase in the error detection rate. Second, in subsequent follow-up analyses (not reported in Staub et al., 2019) we compared eye movements in trials when readers noticed the RT errors, when they failed to notice these errors, and when they read fully grammatical sentences. We found that when readers did not notice the repeated *the*, fixation durations and regression probabilities on each of the two instances of *the* were statistically indistinguishable from eye movements on the single *the* in grammatical sentences. In other words, for the 54% of the trials with a repeated *the* on which participants did not detect this repetition explicitly (i.e., they made a 'no' response to 'Was there anything wrong with that sentence?'), their eye movements revealed no disruption at all, as if they had been reading a grammatical sentence.[5] By contrast, we obtained the expected eye movement disruption on the trials when readers did notice the repetition; they showed inflated reading times on each of the repeated instances of *the*, compared to the single *the* in the grammatical sentences.

Huang and Staub (in press) adopted a similar paradigm to investigate the transposed-word effect reported in Mirault et al. (2018). In each of two experiments, we found a very similar error rate to Mirault et al. (2018), with a transposition of words in an otherwise grammatical sentence going unnoticed around 14% of the time.[6] Note the sizable difference in error rates for repetition, omission, and transposition errors, in experiments that are extremely similar in their design: A repeated *the* was noticed only 46% of the time, an omitted *the* was noticed 67% of the time, a transposition was noticed about 86% of the time, and a repeated noun was noticed about 90% of the time. Following Staub et al. (2019), analyses contingent on fixation patterns showed that directly fixating both transposed words increased the likelihood of detecting a transposition. Again, we also compared eye movements when readers failed to notice a transposition, when they did notice the transposition, and when the sentence was grammatical; the dependent variables were fixation durations (first fixation duration, gaze duration, go-past time and total time) on, and regression probabilities from or back to, the two transposed words. In all six measures, when participants did notice the transposition, their eye movements were significantly disrupted at the second of the two words—the position at which a transposed-word sentence became ungrammatical, in our stimuli—but no statistically significant disruption could be detected when participants did not notice the transposition.[7]

In both of the studies we have just discussed, we also explored the relationship between subjects' reading rate, as measured by their reading time on the filler sentences, and their error detection accuracy in the critical sentences. In Staub et al. (2019), there was a marginally significant correlation ($r = 0.26$) between a subject's mean total sentence reading time on the fillers and detection accuracy in the RT sentences, with slower readers being more accurate; this relationship did not hold for the OT condition. In Huang and Staub (in press), the same relationship reached significance for detection of transposition errors ($r = 0.34$). This pattern might be taken to suggest that failure to notice errors is due in part to fast or inattentive reading. Note, however, that based on the within-subject analyses we have described above, readers do not seem to show particularly fast reading on trials when they fail to notice an error; rather, they seem to read these sentences normally, showing neither disrupted nor speeded reading. Thus, the within-subject analyses do not support the contention that failure to notice errors is due to skimming or inattentive reading.

We have also investigated subject-level factors by comparing different groups of readers. In a small, exploratory eyetracking study, the critical stimuli in Staub et al. (2019) were read by 7- and 8-year-old beginning readers ($N = 17$). While these children showed high comprehension accuracy on filler sentences (92%), and like adults were fairly accurate in accepting the grammatical control sentences (83%) and rejecting the RN sentences (70%), they showed extremely poor accuracy—much worse than adults—for OT sentences (31%) and RT sentences (18%). Thus, child readers were actually completely unable to detect a repeated *the*; they reported an error in RT sentences at essentially the same rate as in fully grammatical sentences where no error was present. However, consistent with many other studies of childrens' reading (Blythe & Joseph, 2011), these beginning readers read very slowly and skipped the critical words only rarely, compared to adults, including fixating quite frequently on both instances of *the* in the RT sentences. Thus, while it may be the case that slower adult readers are slightly more likely to detect errors in the text, this relationship does not extend to children, who are both very slow and very poor at detecting errors.

6 | INTEGRATING RATIONAL INFERENCE WITH EYE MOVEMENTS IN READING

We interpret our eye movement data as consistent in several respects with a rational inference account of failure to notice errors (Gibson et al., 2013). It is not unexpected, for example, that readers are particularly likely to discount perceptual evidence of a repeated *the* or an omitted *the*, compared to evidence of a repeated noun or a transposition of two words. In Staub et al. (2019) we suggested that readers may attribute the appearance of two instances of *the* to an error in eye movement control, interpreting the two instances as a single instance that they have processed twice, first parafoveally and then foveally; this is clearly a noisy-channel-type explanation. And while it is consistent with a rational inference account that readers sometimes fail to notice errors even when they have maximal perceptual input (i.e., they fixate each of the relevant words), such an account would predict that word skipping should lead to some increase in failure to notice errors, as we have found; in Bayesian inference, when perceptual evidence is more equivocal, as is presumably the case when a word is skipped, priors play a greater role. It is also not unexpected that faster readers should notice fewer errors, as this may be seen as reflecting individual readers' model of the noise in the environment (Ryskin et al., 2018), with fast readers being aware (at some level) that their perception of the words on the page may not be completely reliable. It is also not unexpected that beginning readers should be especially likely to discount apparent errors in the text, as they may attribute perceived errors to their own inexperience.

However, our data present one clear challenge for a rational inference account. We see little indication in the eye movement record that errors are perceived, before they are corrected; there is no detectable disruption on those trials on which the reader ultimately fails to notice an error. On its surface, this finding is difficult to reconcile with the idea that noisy-channel inference is genuinely post-perceptual, happening at some delay from the perception of the error itself. It appears, for example, that if a reader initially represents each instance of a repeated *the*, but then infers that there must actually have been only one instance, this inference must happen before any error signal can impede the forward movement of the eyes.

In laying out a potential solution to this puzzle, we begin by briefly sketching how E-Z Reader 10 (Reichle et al., 2009) has modelled readers' response to difficulty in integrating a word into the developing syntactic or semantic representation of the sentence; other eye movement models have not directly tackled this issue. Under E-Z Reader 10, processing of a word consists of four stages. First, visual feature extraction (V) can take place pre-attentively for multiple words at once, allowing targeting of a saccade based on the shape of an upcoming word. Next, word recognition itself, which occurs only for one word at a time, occurs in two sub-stages, denoted L1 and L2. The first stage, L1, is sometimes referred to as a 'familiarity check,' while the completion of the second stage, L2, reflects full lexical access. One important rationale for dividing word recognition into these two stages is to allow a saccade to be programmed to the next word—which happens when L1 has completed—before the currently fixated word has been completely identified. Finally, after a word is recognized, the reader will attempt to incorporate it into the syntactic and semantic context, in an Integration (I) stage.

As just described, E-Z Reader 10 proposes that the integration process occurs after the word has been fully recognized; the word recognition and integration processes occur in entirely separate stages. The integration process normally runs in the background, concurrently with the process of recognition of the next word. But when integration fails, the eyes tend to make a regression to a previous word. Critically, one way that integration can fail is that integration for word *n* is still ongoing when the identification of word *n*+1 is complete (i.e., both L1 and L2 have finished), and word *n*+1 is therefore ready to be integrated. Thus, syntactic and semantic integration is assumed to be perfectly incremental, as word *n* must be integrated before integration of word *n*+1 can begin (Figure 1a).

Within this framework, certain errors might be overlooked, without resulting in disruption in the eye movement record, if the integration process that occurs for each word incorporates a component of rational inference. Let us assume, for example, that in reading the sentence *The mother gave the candle the daughter* (Gibson et al., 2013), the noun phrase *the candle* is initially integrated as the theme argument of *gave*, with the reader then expecting a prepositional phrase headed by *to*, which would include the recipient argument. When the reader now encounters *the daughter* without a preceding *to*, she can either adopt a 'double object' analysis on which *the candle* is now, implausibly, the recipient, and *the daughter* is the theme, or she can analyse *the daughter* as the (plausible) recipient, inferring that an unseen *to* must actually have been present. If the latter, inference-based analysis can be reached quickly enough, during the normal process of integrating *the daughter*, the forward movement of the eyes would not be disrupted.

But in addition, we suggest that the process of semantic and syntactic integration may not be as perfectly incremental as hypothesized by E-Z Reader; in particular, the sentence processing mechanism may sometimes evaluate integrations of multiple words at the same time (Figure 1b). Consider an example transposition sentence from Huang and Staub (in press): *The boy on sat the school bus*. Under E-Z Reader's assumption of maximally incremental integration, the word *on* is always fully integrated, as the beginning of a prepositional phrase modifying *The boy*, prior to any attempt to integrate the word *sat*. As a result, encountering *sat* should result in disruption, because this word constitutes a clear phrase structure violation given the preceding integration of *on*. But we propose that on some occasions, *on* may not be immediately integrated into the sentence representation. Instead, the processor may enter into a state in which the words *on* and *sat* are both temporarily unintegrated. In this case, rational inference may be brought to bear in favour of *sat on* and against *on sat*.

Critically, this account retains the assumption of serial word recognition, which as we noted above is supported by a range of empirical evidence. However, it relaxes the assumption that eye movement disruption will ensue whenever word *n* is not integrated before word *n*+1 is recognized, and instead posits that two (or perhaps more) words may be temporarily held in an unintegrated state. While we agree that sentence processing is highly incremental (e.g., Christiansen & Chater, 2016), we do not think there is direct evidence that it is as perfectly incremental as proposed by E-Z Reader.[8] In particular, the many studies showing that a syntactic or semantic anomaly rapidly disrupts eye movements (e.g., Frazier & Rayner, 1982; Rayner et al., 2004) do not provide this evidence. These studies show that that on average, there is disruption as soon as an anomaly is encountered; for example, mean fixation durations are longer on an anomalous than a non-anomalous word. But an effect on the mean could be driven by a subset of trials where the anomaly is rapidly detected. If there is substantial variability (both within- and between-subjects,

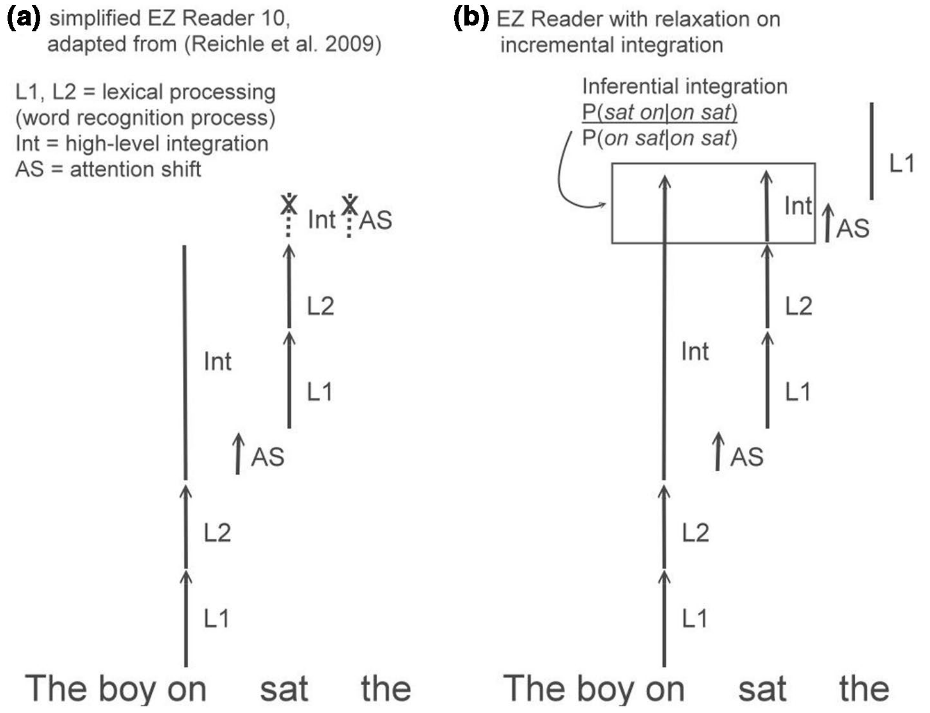

FIGURE 1 (a) Simplified illustration of failure of integration in E-Z Reader 10 (Reichle et al., 2009). If integration of the word *on* is still incomplete (the line without an arrow) when L2 completes for the word *sat*, disruption results and integration of *sat* does not begin. Thus, the words *on* and *sat* are never available for integration at the same time. (b) A revision of the model relaxing the assumption of perfectly incremental integration. A buffer, represented by the rectangle, can hold more than one unintegrated word. If *on* has not been integrated by the time the word *sat* has been recognized, the integration process can consider both possible word orders. Note that the words are still recognized serially and sequentially from left to right; any re-ordering happens at the integration stage

and both within- and between-items) in how rapidly an anomalous word is integrated, we would expect to see effects of anomaly on the mean of various early eye movement measures, but at the same time, variability in the latency of integration could help to explain why readers' eye movements are not always disrupted by an error, especially when rational inference over a multi-word window might support a plausible analysis of the input. As far as we know, no previous research has investigated how syntactic or semantic anomaly influences distributions of fixation durations, as opposed to simply investigating effects on the mean (as has been done for, e.g., the effect of word frequency; Staub et al., 2010); if integration is sometimes delayed, there should be a distinct subset of trials on which effects of anomaly do not show up in first-pass eye movement measures.

7 | METHODOLOGICAL ISSUES

A central empirical result that has informed our theorizing about post-lexical integration and eye movement control is that when an error is not explicitly detected, there is also little or no indication in the eye movement record that the error caused difficulty in the course of incremental processing. A very interesting, and as yet unanswered, question is whether a word repetition or transposition error that is not explicitly detected would leave a signature in other measures, such as EEG. Future research, perhaps involving an eye movement/EEG co-registration paradigm (e.g., Dimigen et al., 2011; Kretzschmar et al., 2015) could investigate this question. We argue that a form of rational inference can take place quickly, and can take place without interrupting the forward progress of the eyes, but it is still entirely possible that such inference is detectable in the EEG record. There are other processing signatures that appear to show up in EEG record but not in eye movements, such as a late positivity associated with integrating an unpredicted word (Frisson et al., 2017; Van Petten & Luka, 2012), and there are other instances in which the dissociation occurs in the opposite direction (e.g., Kretzschmar et al., 2015).

One relevant ERP study is by Wen et al. (2021), who used items adopted from the original transposed-word study (Mirault et al., 2018), and showed a reduced N400 for transposed-word sequences compared to other ungrammatical sequences. However, their analyses only included trials where participants correctly rejected the transposed items. The question of whether a transposition is in some sense registered even when not explicitly detected would be addressed by analysing the complementary set of trials on which subjects did not detect the transposition. Furthermore, Wen et al. adopted a rapid-parallel-visual-presentation design where all words were presented at the same time, and participants had to fixate at the centre without naturally moving their eyes (see Schotter & Payne, 2019, on this issue). On the other hand, Mirault et al. (2020) implemented co-registration of eye-tracking and EEG in a more natural reading paradigm, to probe the parafoveal-on-foveal repetition effect, showing effects on both eye movements and fixation-related potentials that were consistent with the previous literature (Angele et al., 2013; Dare & Shillcock, 2013). However, the boundary paradigm used in this study does not shed light on downstream processes beyond first-pass reading, and does not reveal what is actually encoded by the subject. A future co-registration study might adopt the same paradigm as Staub et al. (2019), which allows both response-based analyses and natural reading.

8 | CONCLUDING REMARKS

We believe that a rational inference account of readers' failure to notice repetition, omission, and transposition errors is extremely promising. At the same time, we have highlighted an apparent paradox: While many eye movement experiments have demonstrated very rapid effects of syntactic or semantic violations in the eye movement record, certain kinds of errors frequently go unnoticed at a conscious level, and also appear to leave no trace in the eye movement record when they are not noticed. We have suggested that the solution to this apparent paradox may involve building rational inference into the post-lexical integration process that follows recognition of each word, and relaxing the assumption that integration is always perfectly incremental; these modifications preserve the assumption that word recognition itself is serial. Clearly, these ideas need to be further fleshed out and computationally implemented, so that their specific predictions can be assessed. At the same time, much more empirical work is needed to identify the conditions under which errors are not noticed; this work is ongoing in our own lab. We suggest that response-contingent analyses of reading time or EEG data are likely to be particularly fruitful in investigating misreading phenomena.

ACKNOWLEDGEMENTS

Thanks to Chuck Clifton for comments on this manuscript. A grant from the U.S. National Science Foundation (BCS 1732008) to AS facilitated this work.

ORCID

Adrian Staub https://orcid.org/0000-0002-9045-9610

ENDNOTES

[1] In this article we restrict ourselves to discussion of readers' failure to notice errors arising from repeated, omitted, or transposed words, mostly neglecting the interesting work on a variety of other 'misreading' phenomena, including bilingual readers' language intrusion errors (e.g., Schotter et al., 2019), the 'missing-letter effect,' (e.g., Koriat et al., 1991) and failure to notice semantic anomalies, as in the so-called 'Moses Illusion' (e.g., Bohan & Sanford, 2008). This is entirely due to space limitation; we do not mean to imply, by excluding these phenomena from the present discussion, that they are not of theoretical relevance.

[2] Because the present focus is on readers' interpretation of the printed text, we focus on Gibson and colleagues' model rather than on rational models of eye movement control itself (Bicknell & Levy, 2010; 2011; Levy et al., 2009; see also Legge et al., 1997). These models emphasize the role of uncertain perceptual evidence and priors in modulating the pattern of readers' eye movements as they inspect a sentence.

[3] Levy's (2008) version of the 'noisy channel' model emphasizes a role for uncertainty in the input itself. On this version, it is not as obvious that noisy channel inference should be regarded as post-perceptual.

[4] We also note that in a recent experiment (Huang & Staub, in prep), we assessed the rate of failure to notice transpositions when sentences were presented in Rapid Serial Visual Presentation (RSVP) format, that is, one word at a time. Remarkably, subjects made about 20% errors. While these are preliminary data, they suggest that failure to notice transpositions can occur even when the order of encoding is fixed—a strong argument that parallel word identification does not play the critical role.

[5] There was some indication of a different pattern for the OT trials; here, there was increased first-pass reading time on the critical noun (swing in (2)) whether or not readers explicitly noticed the error. Potentially, there is a relevant confound: Only in the OT condition was the critical noun not always preceded by a short, high-frequency function word.

[6] Our experiments uncovered differences in error rate based on word length and word class of the transposed words, which we are continuing to investigate, but which for reasons of space we must omit from discussion here.

[7] Bohan and Sanford (2008) carried out a similar response-contingent analysis of readers' failure to notice semantic anomalies, with a similar conclusion, that is, there was disruption in the eye movement record only on those trials when an anomaly was consciously detected.

[8] Reichle et al. (2009) make it clear that the assumption of perfect incrementality of integration is not a theoretical commitment, but serves as a 'placeholder for a deeper theory of postlexical language processing during reading' (p. 6).

REFERENCES

Angele, B., Tran, R., & Rayner, K. (2013). Parafoveal–foveal overlap can facilitate ongoing word identification during reading: Evidence from eye movements. *Journal of Experimental Psychology: Human Perception and Performance*, *39*, 526–538.

Balota, D. A., Pollatsek, A., & Rayner, K. (1985). The interaction of contextual constraints and parafoveal visual information in reading. *Cognitive Psychology*, *17*, 364–390.

Bavelier, D., & Potter, M. C. (1992). Visual and phonological codes in repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 134–147.

Bicknell, K., & Levy, R. (2010, July). A rational model of eye movement control in reading. *In Proceedings of the 48th annual meeting of the Association for Computational Linguistics* (pp. 1168–1178).

Bicknell, K., & Levy, R. (2011). Why readers regress to previous words: A statistical analysis. In Proceedings of the Annual Meeting of the Cognitive Science Society (Vol. 33, No. 33).

Blythe, H. I., & Joseph, H. S. S. L. (2011). Children's eye movements during reading. In S. P. Liversedge, I. D. Gilchrist, & S. Everling (Eds.), *The Oxford handbook of eye movements* (pp. 643–662). Oxford University Press.

Bohan, J., & Sanford, A. (2008). Short article: Semantic anomalies at the borderline of consciousness: An eye-tracking investigation. *Quarterly Journal of Experimental Psychology*, *61*, 232–239.

Brothers, T., Hoversten, L. J., & Traxler, M. J. (2017). Looking back on reading ahead: No evidence for lexical parafoveal-on-foveal effects. *Journal of Memory and Language*, *96*, 9–22.

Brysbaert, M. (2019). How many words do we read per minute? A review and meta-analysis of reading rate. *Journal of Memory and Language*, *109*, 104047.

Brysbaert, M., & Vitu, F. (1998). Word skipping: Implications for theories of eye movement control in reading. In G. Underwood (Ed.), Eye guidance in reading and scene perception (pp. 125–147). Elsevier Science Ltd.

Choi, W., & Gordon, P. C. (2014). Word skipping during sentence reading: Effects of lexicality on parafoveal processing. *Attention, Perception, & Psychophysics*, *76*, 201–213.

Christiansen, M. H., & Chater, N. (2016). The now-or-never bottleneck: A fundamental constraint on language. *Behavioral and Brain Sciences*, *39*, 1–72.

Dare, N., & Shillcock, R. (2013). Serial and parallel processing in reading: Investigating the effects of parafoveal orthographic information on nonisolated word recognition. *Quarterly Journal of Experimental Psychology*, *66*, 487–504.

Dilley, L. C., & Pitt, M. A. (2010). Altering context speech rate can cause words to appear or disappear. *Psychological Science*, *21*, 1664–1670.

Dimigen, O., Sommer, W., Hohlfeld, A., Jacobs, A. M., & Kliegl, R. (2011). Coregistration of eye movements and EEG in natural reading: Analyses and review. *Journal of Experimental Psychology: General*, *140*, 552–572.

Engbert, R., Nuthmann, A., Richter, E. M., & Kliegl, R. (2005). SWIFT: A dynamical model of saccade generation during reading. *Psychological Review*, *112*, 777–813.

Eysenck, M. W., & Keane, M. T. (2005). *Cognitive psychology: A student's handbook* (5th ed.). Taylor & Francis.

Franklin, M. S., Smallwood, J., & Schooler, J. W. (2011). Catching the mind in flight: Using behavioral indices to detect mindless reading in real time. *Psychonomic Bulletin & Review*, *18*, 992–997.

Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, *14*, 178–210.

Frisson, S., Harvey, D. R., & Staub, A. (2017). No prediction error cost in reading: Evidence from eye movements. *Journal of Memory and Language*, *95*, 200–214.

Gibson, E., Bergen, L., & Piantadosi, S. T. (2013). Rational integration of noisy evidence and prior semantic expectations in sentence interpretation. *Proceedings of the National Academy of Sciences*, *110*, 8051–8056.

Goodman, K. S. (1967). Reading: A psycholinguistic guessing game. *Literacy Research and Instruction*, *6*, 126–135.

Harris, C. L., & Morris, A. L. (2000). Orthographic repetition blindness. *The Quarterly Journal of Experimental Psychology Section A*, *53*, 1039–1060.

Healy, A. F., & Zangara, T. K. (2017). Examining misses in reading aloud repeated words. *Quarterly Journal of Experimental Psychology*, *70*, 373–377.

Hemingway, E. (1952). *The old man and the Sea*. Simon & Schuster.

Huang, K.-J., & Staub, A. (in press). Using eye tracking to investigate failure to notice word transpositions in reading. Cognition.

Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, *27*, 117–143.

Kanwisher, N. G., & Potter, M. C. (1990). Repetition blindness: Levels of processing. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 30–47.

Kennedy, A., & Pynte, J. (2005). Parafoveal-on-foveal effects in normal reading. *Vision Research*, *45*, 153–168.

Keshev, M., & Meltzer-Asscher, A. (2021). Noisy is better than rare: Comprehenders compromise subject-verb agreement to form more probable linguistic structures. *Cognitive Psychology*, *124*, 101359.

Kliegl, R., Nuthmann, A., & Engbert, R. (2006). Tracking the mind during reading: The influence of past, present, and future words on fixation durations. *Journal of Experimental Psychology: General*, *135*, 12–35.

Koriat, A., Greenberg, S. N., & Goldshmid, Y. (1991). The missing-letter effect in Hebrew: Word frequency or word function? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 66–80.

Kretzschmar, F., Schlesewsky, M., & Staub, A. (2015). Dissociating word frequency and predictability effects in reading: Evidence from coregistration of eye movements and EEG. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*, 1648.

Legge, G. E., Klitz, T. S., & Tjan, B. S. (1997). Mr. Chips: An ideal-observer model of reading. *Psychological Review*, *104*, 524–553.

Levy, R. (2008). A noisy-channel model of human sentence comprehension under uncertain input. *In Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing* (pp. 234–243).

Levy, R. (2011). Integrating surprisal and uncertain-input models in online sentence comprehension: Formal techniques and empirical results. *In Proceedings of the 49th annual meeting of the Association for Computational Linguistics: Human Language Technologies* (pp. 1055–1065).

Levy, R., Bicknell, K., Slattery, T., & Rayner, K. (2009). Eye movement evidence that readers maintain and act on uncertainty about past linguistic input. *Proceedings of the National Academy of Sciences*, *106*, 21086–21090.

Mirault, J., Snell, J., & Grainger, J. (2018). You that read wrong again! A transposed-word effect in grammaticality judgments. *Psychological Science*, *29*, 1922–1929.

Mirault, J., Yeaton, J., Broqua, F., Dufau, S., Holcomb, P. J., & Grainger, J. (2020). Parafoveal-on-foveal repetition effects in sentence reading: A co-registered eye-tracking and electroencephalogram study. *Psychophysiology*, *57*, e13553.

Morrill, T. H., Dilley, L. C., McAuley, J. D., & Pitt, M. A. (2014). Distal rhythm influences whether or not listeners hear a word in continuous speech: Support for a perceptual grouping hypothesis. *Cognition*, *131*, 69–74.

Ni, W., Fodor, J. D., Crain, S., & Shankweiler, D. (1998). Anomaly detection: Eye movement patterns. *Journal of Psycholinguistic Research*, *27*, 515–539.

Pearlmutter, N. J., Garnsey, S. M., & Bock, K. (1999). Agreement processes in sentence comprehension. *Journal of Memory and Language*, *41*, 427–456.

Poppels, T., & Levy, R. (2016). Structure-sensitive noise inference: Comprehenders expect exchange errors. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Meeting of the Cognitive Science Society* (pp. 378–383). Austin, TX: Cognitive Science Society.

Rayner, K. (1975). The perceptual span and peripheral cues in reading. *Cognitive Psychology*, *7*, 65–81.

Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, *124*, 372–422.

Rayner, K., & Duffy, S. A. (1986). Lexical complexity and fixation times in reading: Effects of word frequency, verb complexity, and lexical ambiguity. *Memory & Cognition*, *14*, 191–201.

Rayner, K., Warren, T., Juhasz, B. J., & Liversedge, S. P. (2004). The effect of plausibility on eye movements in reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 1290–1301.

Reichle, E. D., Pollatsek, A., Fisher, D. L., & Rayner, K. (1998). Toward a model of eye movement control in reading. *Psychological Review*, *105*, 125–157.

Reichle, E. D., Reineberg, A. E., & Schooler, J. W. (2010). Eye movements during mindless reading. *Psychological Science*, *21*, 1300–1310.

Reichle, E. D., Warren, T., & McConnell, K. (2009). Using EZ Reader to model the effects of higher level language processing on eye movements during reading. *Psychonomic Bulletin & Review*, *16*, 1–21.

Reilly, R. G., & Radach, R. (2006). Some empirical tests of an interactive activation model of eye movement control in reading. *Cognitive Systems Research*, *7*, 34–55.

Ryskin, R., Futrell, R., Kiran, S., & Gibson, E. (2018). Comprehenders model the nature of noise in the environment. *Cognition*, *181*, 141–150.

Schad, D. J., Nuthmann, A., & Engbert, R. (2012). Your mind wanders weakly, your mind wanders deeply: Objective measures reveal mindless reading at different levels. *Cognition*, *125*, 179–194.

Schotter, E. R., Li, C., & Gollan, T. H. (2019). What reading aloud reveals about speaking: Regressive saccades implicate a failure to monitor, not inattention, in the prevalence of intrusion errors on function words. *Quarterly Journal of Experimental Psychology*, *72*, 2032–2045.

Schotter, E. R., & Payne, B. R. (2019). Eye movements and comprehension are important to reading. *Trends in Cognitive Sciences*, *23*, 811–812.

Snell, J., & Grainger, J. (2017). The sentence superiority effect revisited. *Cognition*, *168*, 217–221.

Snell, J., & Grainger, J. (2019a). Readers are parallel processors. *Trends in Cognitive Sciences*, *23*, 537–546.

Snell, J., & Grainger, J. (2019b). Word position coding in reading is noisy. *Psychonomic Bulletin & Review*, *26*, 609–615.

Snell, J., Meeter, M., & Grainger, J. (2017). Evidence for simultaneous syntactic processing of multiple words during reading. *PloS One*, *12*(3), e0173720.

Snell, J., vanLeipsig, S., Grainger, J., & Meeter, M. (2018). OB1-reader: A model of word recognition and eye movements in text reading. *Psychological Review*, *125*, 969–984.

Staub, A., Dodge, S., & Cohen, A. (2019). Failure to notice function word repetitions and omissions in reading: Are eye movements to blame? *Psychonomic Bulletin & Review*, *26*, 340–346.

Staub, A., White, S. J., Drieghe, D., Hollway, E. C., & Rayner, K. (2010). Distributional effects of word frequency on eye fixation durations. *Journal of Experimental Psychology: Human Perception and Performance*, *36*, 1280–1293.

Van Petten, C., & Luka, B. J. (2012). Prediction during language comprehension: Benefits, costs, and ERP components. *International Journal of Psychophysiology*, *83*, 176–190.

Wen, Y., Mirault, J., & Grainger, J. (2021). Fast syntax in the brain: Electrophysiological evidence from the rapid parallel visual presentation paradigm (RPVP). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *47*, 99–112.

White, A. L., Palmer, J., & Boynton, G. M. (2018). Evidence of serial processing in visual word recognition. *Psychological Science*, *29*, 1062–1071.

White, A. L., Palmer, J., & Boynton, G. M. (2020). Visual word recognition: Evidence for a serial bottleneck in lexical access. *Attention, Perception, & Psychophysics*, *82*, 2000–2017.

White, A. L., Palmer, J., Boynton, G. M., & Yeatman, J. D. (2019). Parallel spatial channels converge at a bottleneck in anterior word-selective cortex. *Proceedings of the National Academy of Sciences*, *116*, 10087–10096.

AUTHOR BIOGRAPHIES

Kuan-Jung Hung is a PhD student in the Department of Psychological and Brain Sciences at the University of Massachusetts Amherst. He came to UMass after obtaining a degree in Foreign Languages and Literature at National Cheng Kung University in Taiwan. He is interested in word recognition and syntactic processes, and their interplay, in sentence processing, including both English and Mandarin Chinese.

Adrian Staub is a Professor in the Department of Psychological and Brain Sciences at the University of Massachusetts Amherst. His research focuses on understanding the details of how readers recognize printed words and analyse sentence structure, primarily by means of experiments in which readers' eye movements are monitored. He is the Director of the UMass Eyetracking Lab.

How to cite this article: Huang, K.-J., & Staub, A. (2021). Why do readers fail to notice word transpositions, omissions, and repetitions? A review of recent evidence and theory. *Language Linguistics Compass*, e12434. https://doi.org/10.1111/lnc3.12434