**Estimating the true cost of garden-pathing: A computational model of
latent cognitive processes**

Dario Paape and Shravan Vasishth
Department of Linguistics, University of Potsdam

Author Note

Please send correspondence concerning this article to: Dario Paape, Department
Linguistik, Universität Potsdam, Karl-Liebknecht-Straße 24–25, D-14476, Germany.
Email: paape@uni-potsdam.de.

Abstract

What is the processing cost of being garden-pathed by a temporary syntactic ambiguity? We argue that comparing average reading times in garden-path versus non-garden-path sentences is not enough to answer this question. Trial-level contaminants such as inattention, the fact that garden-pathing may occur non-deterministically in the ambiguous condition, and “triage” (rejecting the sentence without reanalysis; Fodor & Inoue, 2000) lead to systematic underestimates of the true cost of garden-pathing. Furthermore, the pure garden-path effect due to encountering an unexpected word needs to be separated from the additional cost of syntactic reanalysis. To get more realistic estimates for the individual processing costs of garden-path surprisal and syntactic reanalysis, we implement a novel computational model that includes trial-level contaminants as probabilistically occurring latent cognitive processes. The model shows good predictive fit to existing reading-time and judgment data. Furthermore, the latent-process approach captures differences between NP/Z garden-path sentences and RRC garden-path sentences: The NP/Z garden path occurs nearly deterministically but can be mostly eliminated by adding a comma. By contrast, the RRC garden path occurs with a lower probability, but disambiguation via semantic plausibility is not always effective.

Keywords: garden-path effect, syntactic reanalysis, multinomial processing tree, latent processes, mixture modeling

Estimating the true cost of garden-pathing: A computational model of latent cognitive processes

Introduction

What is the cost of garden-pathing? The most common approach to answering this question is to compare reading times between temporarily ambiguous sentences and control sentences, as in (1a,b):¹

- (1) (a) The news show stated that while the team trained, the striker wondered whether the damage would take long to heal.
- (b) The news show stated that while the team trained the striker wondered whether the damage would take long to heal.

In (1b), the human sentence processor has a strong tendency to attach *the striker* as the direct object of the verb *trained* on the first pass. This analysis turns out to be incorrect when the verb *wondered* is encountered, and has to be abandoned in favor of an analysis in which *trained* is used intransitively. By contrast, in the control sentence (1a), the comma after *trained* leads to early disambiguation of the structure, which prevents the garden path. A plausible measure of the cost of garden-pathing would thus be the difference in mean reading times between (1a) and (1b) at the word *wondered*, calculated from a sample of multiple sentences with a structure like (1a,b) and multiple participants. A consistent finding is that *wondered* is read more slowly on average in (1b) than in (1a) (e.g., Grodner, Gibson, Argaman, & Babyonyshev, 2003; Sturt, Pickering, & Crocker, 1999; Van Schijndel & Linzen, 2021).

Despite the relatively simple manipulation, even after many years of research, it is still an open question which latent processes are reflected by the observed slowdown in the garden-path condition. It may be any combination of the following: the initial surprise when encountering the word *wondered* at an unexpected point in the sentence, the search for the grammatical subject of *wondered*, the surprise when no available subject can be found nearby, the search for an alternative syntactic analysis,² and finally the suppression or deconstruction of the initial analysis and the building of a new syntactic structure (see Meng & Bader, 2000 and Van Schijndel & Linzen, 2021 for discussion).

In addition to the theoretical treatment of garden-path effects, there is also a question of whether the observed difference in reading times between (1a) and (1b) is really an accurate measure of the processing cost incurred by garden-pathing. Classical garden-path theory assumes that garden-pathing is deterministic: Any sentence like (1b) should always be more difficult to process than a matched control sentence like (1a), because the human sentence processor invariably prefers immediate local attachment (Frazier, 1979). However, more recent models of syntactic ambiguity resolution assume that garden-pathing is non-deterministic (e.g., Logachev & Vasisht, 2016; Van Gompel, Pickering, Pearson, & Liversedge, 2005; Van Gompel, Pickering, & Traxler, 2000), that is, it may sometimes occur and sometimes not occur. Non-deterministic models assume that

¹ The sentences are taken from Mitchell, Shen, Green, and Hodgson (2008).

² We assume a serial parser for the sake of simplicity. Parallel parsing models are discussed in the final section of the paper.

occasionally, the dispreferred analysis may be adopted instead of the preferred analysis on the first pass, that is, *the striker* may sometimes *not* be attached to *trained* even in the absence of a disambiguating comma. In such trials, there would be no garden path: There would be no surprise at encountering *wondered*, and no syntactic reanalysis would be necessary. If such trials occur in the data, they will pull the mean difference between the garden-path and non-garden-path conditions towards zero, resulting in an underestimate of the true garden-path cost.

Another reason why the cost of garden-pathing may be underestimated is that readers may not always be fully attentive or fully motivated during reading (Franklin, Smallwood, & Schooler, 2011; Smallwood, 2011; Unsworth & McMillan, 2013). In some trials, readers may be inattentive or simply have no incentive to spend time trying to fully understand the sentence (e.g., Ferreira & Yang, 2019; Logačev & Dokudan, 2021; Swets, Desmet, Clifton, & Ferreira, 2008; Zhang, Qu, Miller, & Cortina, 2020). A completely inattentive reader may skip the word *wondered* in (1b) or look at it without processing it deeply, showing no difference between (1a) and (1b). A somewhat more attentive but unmotivated reader may notice the unexpectedness of *wondered* in (1b) and show a small slowdown in reading, but fail to carry out syntactic reanalysis. Instead, the unmotivated reader may conclude that the sentence is ungrammatical or use pragmatic reasoning to guess a likely interpretation (e.g., Christianson, Hollingworth, Halliwell, & Ferreira, 2001). As in the case of non-determinism, aggregating over attentive/motivated trials and inattentive/unmotivated trials results in an underestimate of the garden-path effect that does not reflect the underlying complex reality.

One step towards obtaining more realistic estimates for effects of theoretical interest is to explicitly model “contaminating” sources of trial-level variability such as guessing or lapses in attention (e.g., McVay & Kane, 2012; Ollman, 1966). In the study of reaction times, one can assume that the observations follow a mixture distribution (Ratcliff & Tuerlinckx, 2002): Some percentage of trials reflect the process of interest – say, garden-pathing – while the remaining trials are generated from one or multiple contaminant distributions. Using modern statistical techniques, the percentage of trials belonging to each mixture component, as well as the parameters of each distribution – means and standard deviations – can be estimated from the data (e.g., Vasisht, Jäger, & Nicenboim, 2017). By moving away from the assumption that overall mean reaction times directly reflect cognitive processes of interest, much can be learned about the true generating mechanisms behind the data.

In what follows, we implement a computational model of garden-pathing and syntactic reanalysis that accounts for effects of inattention and non-determinism. The model is informed by the distribution of reading times at the disambiguating region, and also links the reading times to behavioral measures on a trial-by-trial basis. In particular, end-of-trial acceptability judgments and rereading of earlier sentence regions are used to tease apart the latent processes that occur during reading. Garden-pathing in sentences like (1b) is known to cause regressions and rereading in eye-tracking studies, which may be connected to syntactic reanalysis (e.g., Frazier & Rayner, 1982; Mitchell et al., 2008), as well as leading to negative acceptability judgments, which may signal failure to compute a syntactic structure (Ferreira & Henderson, 1991; McElree, 1993; Tabor & Hutchins, 2004; Van Dyke & Lewis, 2003).

We model complete inattention as a source of trials in which garden-pathing is completely absent, as well as cases in which readers are attentive but do not attempt reanalysis at the disambiguating region and instead reject the sentence as ungrammatical. This strategy has been called “trriage” (Fodor & Inoue, 2000; Meng & Bader, 2000). Moreover, our model takes into account the distinction between so-called covert and overt reanalysis: Overt reanalysis involves rereading of earlier parts of the sentence, while covert reanalysis is carried out “in place” at the disambiguating region (Lewis, 1998). Finally, the model assumes that even when reanalysis is attempted it is not always successful (e.g., Ferreira & Henderson, 1991), with the probability of success depending, *inter alia*, on the specific syntactic construction.

A latent-processes model of garden-pathing and reanalysis

The model takes the form of a multinomial processing tree (MPT; Batchelder & Riefer, 1999; Stahl, 2006; Ulrich, 2009). The model structure is shown in Figure 1. For other recent work that applies MPTs to sentence processing, see Logačev and Dokudan (2021) and Paape, Avetisyan, Lago, and Vasisht (2021). Our Bayesian MPT model is implemented in Stan (Stan Development Team, 2021) and fitted via the *rstan* package (Stan Development Team, 2020) in R (R Core Team, 2021). We fit the model to the experimental data of Paape and Vasisht (2021), which contains reading measures – including regressions – and acceptability judgments from 193 participants for two garden-path constructions (see Experimental data section below). To inform the model parameters, trials are categorized into four groups according to two factors: The judgment given after reading the sentence (accept/reject) and the presence or absence of a first-pass regression from the disambiguating region. In combination with the reading times at the disambiguating region, the four trial categories allow inferences about the latent processes that took place in a given trial. The two main linking assumptions are that positive acceptability judgments signal the computation of the correct syntactic structure, and that regressions from the disambiguating region signal overt reanalysis.³

The main sources of information in an MPT are the relative proportions of response categories in the observed data. In an MPT, cognitive processes are assumed to occur serially with some to-be-estimated probability and some to-be-estimated processing cost, with each process feeding into the next one, and ultimately resulting in an observed outcome. A given outcome may be reached via different paths through the tree. For instance, in the model shown in Figure 1, trials in which the subject makes *no* regression at the disambiguating region and ends up rejecting the sentence can be generated in three ways: Inattention during reading, followed by a biased guess of the sentence’s acceptability, garden-pathing followed by triage, that is, the decision not to carry out reanalysis, and finally garden-pathing followed by unsuccessful covert (in-place)

³ These assumptions are necessary simplifications. Negative acceptability judgments could, in principle, also arise due to high processing complexity (Hofmeister, Jaeger, Arnon, Sag, & Snider, 2013), while regressions may sometimes be made simply to verify earlier input (see discussion in Paape & Vasisht, 2021). Participants in the study of Paape and Vasisht (2021) were instructed to reject sentences if they violated the rules of English grammar, or if they were contradictory or nonsensical, which should make rejections due to complexity alone highly unlikely. Furthermore, we assume that regressions and positive acceptability judgments are a better indicator of reanalysis than, for instance, comprehension questions, whose responses may be affected by factors unrelated to online sentence processing (Bader & Meng, 2018).

Figure 1

Multinomial processing tree model of garden-pathing and reanalysis.

reanalysis. As the probability of each path is the product of the probabilities of its component processes and the probability of each outcome is the sum of the relevant path probabilities, the probabilities of individual processes can be inferred from the observed outcome proportions.

Our model is specified hierarchically, that is, the probabilities of individual processes may differ between participants as well as between sentences. The seven probabilities to be estimated from the data are the following:

- $p_{attentive}$: The probability that the participant is reading attentively in a given trial. Varies between participants.
- p_{bias} : The probability of the participant giving an “accept” judgment if they were inattentive when reading the disambiguating region. Varies between participants.
- p_{gp} : The probability of that the participant is garden-pathed, that is, is surprised by the disambiguating word. Varies between participants, sentences, conditions, and syntactic constructions.

- *p_{trriage}*: The probability of triage, that is, of rejecting the sentence without attempting reanalysis. Varies between participants, sentences, and syntactic constructions.
- *p_{covert}*: The probability of carrying out reanalysis covertly rather than overtly. Varies between participants, sentences, and syntactic constructions.
- *p_{covert_success}*: The probability of success of a covert reanalysis attempt. Varies between participants, sentences, and syntactic constructions.
- *p_{overt_success}*: The probability of success of an overt reanalysis attempt. Varies between participants, sentences, and syntactic constructions.

Conditional on the trial category (accept/no regression, reject/no regression, accept/regression, reject/regression), first-pass reading times at the disambiguating region are assumed to reflect the summed cost of the individual processes along the relevant paths through the tree. Separate costs are assumed for paying attention, for being garden-pathed – that is, for registering that the disambiguating word is unexpected – for attempting covert reanalysis, and for making a regression. The processing cost of reanalysis is assumed to be the same whether the attempt is successful or not. Given that the same outcomes can be generated via different paths with different costs, the resulting reading times follow a mixture distribution (Heck & Erdfelder, 2016; Klauer & Kellen, 2018; McLachlan & Peel, 2000). Processing costs are not assumed to vary between participants or sentences. Assuming variation in both process probabilities and process costs would lead to an unidentifiable model, as it would not be possible to distinguish between relatively cheap processes that occur often and relatively costly processes that only occur rarely.

Depending on whether one or multiple paths can lead to an outcome, reading times in the model are assumed to follow either a lognormal distribution with mean μ and standard deviation σ (one path) or a mixture of multiple lognormal distributions with different μ and a shared σ (multiple paths). All estimated costs in the model are constrained to be positive and are added to a baseline μ . Another assumption encoded in the model is that reading times in inattentive trials are less variable than reading times in attentive trials. Research on mindless reading has shown that readers are less sensitive to factors like word length and word frequency when reading mindlessly (Bixler & D’Mello, 2015; Franklin et al., 2011; Reichle, Reineberg, & Schooler, 2010; Schad, Nuthmann, & Engbert, 2012). We thus assume that reading times should have a smaller σ in inattentive trials than in attentive trials, and constrain the parameters accordingly.

Experimental data

The study of Paape and Vasishth (2021) is especially interesting for our purposes because it compared two types of garden-path constructions within the same large participant sample (193 subjects): Sentences with an NP/zero complement (NP/Z) ambiguity as discussed in the introduction and sentences with a main clause/reduced relative clause (RRC) ambiguity, as shown in Table 1. The two sentence types differ in terms their mode of disambiguation (orthographic versus semantic) and in terms of the magnitude of the average garden-path effect (Van Schijndel & Linzen, 2021), which should affect the latent processes in our MPT model.

NP/Z AMBIGUITY: Comma (unambiguous) versus no comma (ambiguous)

- (2) The news show · stated · that · while · the team · $\left\{ \begin{array}{l} \text{trained,} \\ \text{trained} \end{array} \right\}$ · the striker ·
wondered · whether · the damage · would take long to heal.

RRC AMBIGUITY: Inanimate NP (unambiguous) versus animate NP (ambiguous)

- (3) Things · soon · turned ugly, · because · $\left\{ \begin{array}{l} \text{the package} \\ \text{the lawyer} \end{array} \right\}$ · sent · **by the**
governor · was neglected · by the officials · in the town hall.

Table 1

Example stimuli used in the experiment of Paape and Vasishth (2021). The disambiguating region is bolded. The · symbol marks region boundaries.

In the RRC ambiguity, the NP *the package/the lawyer* can initially be attached either as the agent of *sent*, until the PP *by the governor* arrives to signal a passive structure (*the package (that was) sent . . .*). In contrast with NP/Z sentences, the “control” condition in RRC sentences is disambiguated via semantic plausibility: A package is not a plausible agent of a sending event while a lawyer is.

There has been some discussion in the literature about whether semantic disambiguation can prevent garden-path effects or not, that is, about whether first-pass parsing decisions in RRC sentences are influenced by plausibility (Clifton et al., 2003; Ferreira & Clifton, 1986; Trueswell, Tanenhaus, & Garnsey, 1994). In the context of the present work, the aim is not to give a categorical “yes” or “no” answer to this question, but to get an estimate of *how* well semantic disambiguation works: A unique advantage of the MPT model is that it can estimate probabilities of garden-pathing for both unambiguous and ambiguous versions of NP/Z and RRC sentences. A high probability of garden-pathing even in the “unambiguous” RRC condition would signal that semantic disambiguation is not always effective.

The data of Paape and Vasishth (2021) were collected using a web-based bidirectional self-paced reading (BSPR) paradigm. The BSPR paradigm extends regular moving-window self-paced reading by giving participants the opportunity to move backward in addition to forward through the sentence in a region-by-region fashion (Paape & Vasishth, 2022). The right arrow key reveals the next word or region (and masks the previous one) while the left arrow key reveals the previous word (and masks the next one). The time between consecutive key presses is the reading time for the region in question. Participants can also return directly to the beginning of the sentence by pressing the ESC key or jump directly to the acceptability judgment by pressing the CTRL key.

In BSPR, reading time measures can be computed similarly to eye tracking. For instance, in the partial example trial shown in (4), the first-pass reading time for the word *wondered* would be the time spent on this word prior to making a regression, and later rereading would be excluded from this measure.

(4) trained --- ----- ----- ----- Key press: →
 ----- the striker ----- ----- Key press: →
 ----- ----- wondered ----- Key press: ← (first-pass regression)
 ----- the striker ----- ----- Key press: ← (second-pass regression)
 trained --- ----- -----

Figure 2 summarizes the relevant dependent variables for each sentence type (NP/Z vs RRC) in the Paape and Vasishth (2021) data: First-pass reading times at the disambiguating region, first-pass regressions, and end-of-sentence acceptability judgments, as well as their relationships.

Figure 2

Top: First-pass reading times at the disambiguating region by sentence type, regression (present or absent), and condition. Middle: Regression proportion by sentence type and condition. Bottom: Proportion of positive acceptability judgments by sentence type, regression, and condition. Error bars show 95% confidence intervals.

As the figure shows, the average garden-path effect in first-pass reading times and regression probabilities is larger in the NP/Z construction than in the RRC construction. Fewer positive acceptability judgments are given for ambiguous sentences and, especially in the NP/Z construction, when the trial contains a regression from the disambiguating region. Across both constructions, first-pass reading times are longer when they are followed by a regression. This contrasts with eye tracking, where the opposite pattern has been observed (Altmann, Garnham, & Dennis, 1992; Rayner & Sereno, 1994), and could be due to the more deliberate nature of regressions in the BSPR paradigm.

We now describe how the MPT model was fitted to the data.

Model fitting and results

Data points with very short first-pass reading times (below 150 ms) or very long first-pass reading times (above 5000 ms) were removed prior to fitting the model, resulting in a loss of 2% of the data. For the remaining trials, the factor condition was sum-coded (unambiguous -1 , ambiguous $+1$), as was the factor construction (RRC -1 , NP/Z $+1$). Four MCMC chains were run with 4000 iterations each, the first 2000 of which were discarded as warm-up. We report percentile-based 95% credible intervals for all parameters of interest. For better interpretability, parameters on the log scale have been back-transformed to the raw millisecond scale.

The estimated costs and process probabilities from the model are as follows:

- **Attention probability** ($p_{attentive}$): Participants are attentive in 68–79% of trials.
- **Judgment bias** (p_{bias}): Participants are highly biased towards judging a sentence as acceptable when they didn't read the disambiguating region attentively (75–87% acceptable).
- **Garden-path probability** (p_{gp}): Garden-pathing is more likely in ambiguous NP/Z sentences (99% garden-pathing) than in ambiguous RRC sentences (65–87% garden-pathing). Disambiguation by a comma in NP/Z sentences is highly effective at preventing garden-pathing (3–11% garden-pathing) compared to semantic disambiguation in RRC sentences (35–65% garden-pathing).
- **Triage probability** ($p_{trriage}$): Triage occurs more often in NP/Z sentences (42–66% triage) compared to RRC sentences (3–10% triage).
- **Covert reanalysis probability** (p_{covert}): Reanalysis is less likely to be covert in NP/Z sentences (66–80% covert) than in RRC sentences (89–94% covert).
- **Covert reanalysis success probability** ($p_{covert_success}$): Covert reanalysis succeeds less often in NP/Z sentences (32–54% successes) than in RRC sentences (75–88% successes).
- **Overt reanalysis success probability** ($p_{overt_success}$): Overt reanalysis also succeeds less often in NP/Z sentences (20–44% successes) than in RRC sentences (67–88% successes).
- **Attention cost**: Being attentive as opposed to inattentive results in a slowdown of 28–68 ms.
- **Garden-path cost**: Being garden-pathed by the unexpected disambiguating word(s) costs 35–90 ms.
- **Reanalysis cost**: Covert reanalysis costs 579–690 ms.

- **Regression cost:** Making a regression from the disambiguating region costs 494–617 ms.

The parameter estimates are plausible given what is known about the constructions. However, it is also important to check if the model can generate data that look similar to the observed experimental data, especially given the large estimated costs for reanalysis and regressions. To this end, we perform graphical posterior predictive checks (Gelman et al., 2014). Data are repeatedly generated from the fitted model to obtain the posterior predictive distribution, and the generated data are compared to the observed data. The higher the overlap between the observed data and the generated data, the better the model can be assumed to generalize to future data.

Figure 3 shows graphical posterior predictive checks for the different trial types (accept versus reject, regression versus no regression from disambiguating region). Figure 4 shows graphical posterior predictive checks for log first-pass reading times at the disambiguating region. As the plots show, the model is able to reproduce both the qualitative and the quantitative patterns across both syntactic constructions.

Figure 3

Graphical posterior predictive checks of outcome proportions. Diamonds represent empirical means, black bars represent mean ± 2 SD, violins represent data generated by the model (2000 samples).

The posterior predictive checks are encouraging, but there is also a question of how much model complexity is needed to account for the patterns in the data. If a simpler model is able to predict the data as well as a more complex model, the simple model with fewer assumptions should be preferred.

In order to investigate the contribution of specific latent processes in the MPT to the model's predictive fit, we carry out leave-one-out cross-validation as implemented in the loo package (Vehtari, Gabry, Yao, & Gelman, 2019). True leave-one-out cross-validation would fit the model n times, n being the number of data points, holding

Figure 4

Graphical posterior predictive checks of log first-pass reading times at the disambiguating region. Diamonds represent empirical means, black bars represent mean ± 2 SD, violins represent data generated by the model (1000 samples).

out exactly one data point each time. The expected log predictive densities ($\widehat{elpd}$) of all n held-out data points are then summed, yielding a measure of how well the model predicts unseen data. The loo package uses a more efficient approximation of this process called PSIS-LOO (Vehtari, Gelman, & Gabry, 2015). Models with too few parameters will not predict unseen data well, because they fail to pick up on structure in the fitted data (underfitting). On the other hand, models with too many parameters will *also* not predict unseen data well, because they become too heavily attuned to the fitted data (overfitting).

The latent processes whose contribution we are most interested in are inattention and triage, as these processes presumably result in underestimates of garden-path effects. We thus compare the following models:

- *full*: The full MPT model as shown in Figure 1.
- *same-triage*: The *full* model with no difference in triage probabilities between NP/Z and RRC sentences.
- *no-triage*: The *full* model without triage.
- *minimal*: The *full* model without inattention and without triage.

The $\widehat{elpd}$ values for each of the models are shown in Table 2. Higher $\widehat{elpd}$ means

that the held-out data are more likely under a given model, indicating better predictive fit. Note that the $\widehat{elpd}$ values in Table 2 are negative, meaning that smaller absolute values indicate better predictive fit.

model	$\widehat{elpd}$	SE
full	−62574	145
same-triage	−62591	145
no-triage	−62811	143
minimal	−63249	141

Table 2

$\widehat{elpd}$ and standard errors by model. The highest $\widehat{elpd}$ value is shown in boldface.

In order to get a sense of whether the differences between models are meaningful, we compute the standard error of the difference in $\widehat{elpd}$ (Vehtari, Gelman, & Gabry, 2017). If the absolute difference is greater than 2 times its associated standard error, one model offers a decisively better fit than the other. Given that $\widehat{elpd}$ descriptively increases with model complexity in Table 2, we make pairwise comparisons between each adjacent pair of less versus more complex models. The results are as follows:

- The *full* model outperforms the *same-triage* model ($\widehat{\Delta elpd} = 18$, SE = 5).
- The *same-triage* model outperforms the *no-triage* model ($\widehat{\Delta elpd} = 220$, SE = 29).
- The *no-triage* model outperforms the *minimal* model ($\widehat{\Delta elpd} = 438$, SE = 42).

Overall, model complexity decisively increases predictive fit. The largest advantage comes from the assumption of occasional inattention in the *no-triage* model compared to the *minimal* model. Adding the triage process in the *same-triage* model further increases predictive fit. Finally, assuming different probabilities of triage in NP/Z versus RRC sentences leads to another small but decisive increase in predictive fit for the *full* model over the *same-triage* model.

Discussion

The cross-validation results show that the best predictive fit to the data is provided by the full MPT model including inattention as a latent process and different triage probabilities for NP/Z and RRC sentences. This suggests that models that do not take these mechanisms into account are not sufficient to explain the empirical patterns. According to the estimate from the full model, participants are inattentive in about one-fifth to one-third of trials. Out of the remaining attentive trials, when garden-pathing occurs in NP/Z sentences, triage is applied in around one-third to half of trials, meaning that no reanalysis is attempted and the sentence is rejected. Compared to NP/Z sentences, triage is much less likely in RRC sentences, but still occurs with non-zero probability.

Several of the estimates from the model point towards NP/Z garden paths being more difficult than RRC garden paths, consistent with the observed reading measures of

Paape and Vasishth (2021) and other recent experimental results (Van Schijndel & Linzen, 2021). According to the model, the garden path in NP/Z sentences is close to being deterministic: In NP/Z sentences without commas, readers are surprised by the main verb in 99% of trials. Deterministic garden-pathing would be in line with the original garden-path theory of Frazier (1979); however, ambiguous RRC sentences do not show a fully deterministic pattern. One possible explanation is that the parser always adopts the structure that is computed most quickly, as assumed by the Unrestricted Race Model (Van Gompel et al., 2005, 2000). In NP/Z sentences, the processing speed advantage of the NP-complement structure over the zero-complement structure may be so large that the race between the syntactic alternatives is essentially pre-decided. By contrast, the speed difference between the main clause and reduced relative clause analyses in RRC sentences may not be as large, resulting in more variability with regard to garden-pathing.

Having factored out the contaminating processes of inattention and triage, as well as non-determinism in parsing, the MPT model also yielded estimates for the processing costs of surprisal-based garden-pathing and reanalysis. The estimated cost of garden-pathing (35–90 ms) is in line with estimates from the previous literature in which differences between condition means were reported, which however vary widely (Van Schijndel & Linzen, 2018).⁴ After having been garden-pathed, the parser can either apply triage or carry out covert reanalysis, which is very costly (579–690 ms). A third alternative, namely making a regression, is also very costly (494–617 ms), which however may be due to the more deliberate nature of regressive rereading in the BSPR paradigm compared to natural reading (Paape & Vasishth, 2022).

In addition to their high cost, both covert and overt reanalysis are very likely to fail in NP/Z sentences. This is consistent with previous findings from comprehension tasks (see Fujita & Cummings, 2021 for a review), and with the observation of Paape and Vasishth (2021) that regressions in NP/Z sentences were associated with fewer rather than more positive acceptability judgments. In RRC sentences, by contrast, reanalysis is often, if not always, successful. We discuss possible reasons for this difference in the final section of the paper.

General discussion

We set out to get an estimate of the true cost of garden-pathing and syntactic reanalysis. To this end, we took into account trial-level contaminants in the reading time measures due to latent cognitive processes: Occasional inattention on part of participants, trials in which garden-pathing is absent due to non-determinism in parsing, and trials in which “triage” is applied and no reanalysis is attempted (Fodor & Inoue, 2000; Meng & Bader, 2000). Our modeling results show that these contaminations make up a sizable amount of trials, with the much-researched NP/Z ambiguity being especially prone to triage.

⁴ The NP/Z effect in Grodner et al. (2003) with 35 participants was 53–78 ms, but the NP/Z effect in Sturt et al. (1999) with 24 participants was 86–218 ms. The NP/Z effect observed by Prasad and Linzen (2019) with 267 participants was 19–37 ms. By comparison, the estimated mean difference for NP/Z sentences in the data of Paape and Vasishth (2022) is 151–209 ms, which may be partly driven by the judgment task.

Assuming a triage mechanism and taking into account end-of-trial acceptability judgments allowed us to separate garden-pathing from syntactic reanalysis: Trials with triage are faster than trials with reanalysis and lead to negative judgments, but they are slower than trials in which garden-pathing is entirely absent and the judgment is positive. Our model showed an increase in processing times between triage trials and non-garden-path trials, which is the pure cost of garden-pathing without reanalysis. This cost is probably due to surprisal, that is, the unexpectedness of the disambiguating word given the previous words in the sentence, which has been independently argued to have a strong influence on reading times (Hale, 2001; Levy, 2008).

However, surprisal is not the full picture. Our model estimates show that reanalysis adds a substantial cost on top of the surprisal cost, in line with recent findings by Van Schijndel and Linzen (2021). Van Schijndel and Linzen calculated surprisal values – that is, negative log probabilities – for disambiguating words in different types of garden-path sentences and used these to predict disambiguation costs in human self-paced reading data. They found that the surprisal model could mostly account for the qualitative patterns in the data but severely underestimated the magnitude of the cost incurred by disambiguation. Their conclusion is that additional mechanisms beyond surprisal, such as syntactic reanalysis, are needed to account for the observed differences in reading times in the data. According to our model estimate, covert syntactic reanalysis takes more than half a second, which is very costly by psycholinguistic standards, where effect estimates are typically below 100 ms.⁵

With regard to the component mechanisms involved in this costly process of reanalysis, our model is compatible with multiple theories. For instance, it is possible that the difficulty of reanalysis lies mainly in identifying the parts of the incorrect initial analysis that need to be changed (e.g., Fodor & Inoue, 1994). Other proposals that our model is compatible with include costs for prosodic restructuring (e.g., Grodner et al., 2003) and for restructuring the syntactic tree (e.g., Sturt, 1996; Sturt et al., 1999). Alternatively, competition-based models (e.g., Spivey & Tanenhaus, 1998) assume that two syntactic analyses are computed in parallel, and that disambiguation towards the dispreferred analysis leads to a costly “reranking” of the structures. In our model, we have assumed that only one analysis is chosen on the first pass, but it is possible that the alternative analysis is computed and then “set aside” until it is needed for reanalysis (Gibson, 1991).

One important aspect of our results that can inform the theory of reanalysis are the observed differences between NP/Z and RRC garden paths. According to our model estimates, garden-pathing is almost guaranteed in NP/Z sentences, as predicted by the deterministic parsing principle Late Closure (attach to the most recent phrase; Frazier, 1979). A fully deterministic model of parsing, however, cannot account for the occasional absence of garden-pathing in RRC sentences. We believe that the Unrestricted Race Model of Van Gompel et al. (2005, 2000) is best suited to explain the difference between the two garden-path types: If one assumes that the processing speed advantage of the preferred structure over the dispreferred structure is much larger in NP/Z sentences than

⁵ This is true for a range of phenomena, including surprisal (Vasishth, Mertzen, Jäger, & Gelman, 2018), similarity-based interference (Jäger, Engelmann, & Vasishth, 2017), and relative clauses (Nicenboim, Schad, & Vasishth, 2021, Chapter 6).

in RRC sentences, the result will be (near) determinism in the former but non-determinism in the latter.

However, NP/Z sentences do not only differ from RRC sentences with regard to the garden-path probability in the ambiguous condition. Disambiguated NP/Z sentences with commas also showed a much lower probability of garden-pathing than disambiguated RRC sentences with animate NPs. This finding is in line with claims by Ferreira and Clifton (1986) and Clifton et al. (2003) that RRC ambiguities still incur processing costs even in the presence of semantic “disambiguation”, suggesting that syntactic preferences can override semantics. At the same time, our estimates show that semantic disambiguation *can* be successful, suggesting that the syntactic module is not completely insulated from semantic influences on the first pass (Trueswell et al., 1994). This suggests a solution to the debate between proponents of “syntax-first” and multiple-constraint accounts of ambiguity resolution: In a non-deterministic parsing model like the Unrestricted Race Model, syntactic preferences may have a stronger effect on the relative processing speed of the competing analyses than semantic constraints. Given the probabilistic effect of semantic cues on parsing decisions, comparing average reading times across conditions may not be sufficient to reliably detect semantic effects. This conclusion underlines the value of implemented cognitive models in adjudicating between psycholinguistic theories: Manipulations that may not always measurably affect average reading times may nevertheless have discernible effects on latent cognitive processes.

A final difference between NP/Z sentences and RRC sentences emerged with regard to the probability of triage and the probability of successful reanalysis. Different probabilities of triage for different sentence types are compatible with a powerful triage mechanism: Triage appears to have insight into the costs and the relative probabilities of success of potential reanalysis attempts, which would need to be learned from experience (Fodor & Inoue, 2000). If reanalysis is judged to be too difficult in a particular trial, none is performed. The estimated probabilities of successful reanalysis in NP/Z versus RRC sentences from the model show that NP/Z sentences are indeed much harder to reanalyze than RRC sentences, making triage attractive for the reader.⁶ NP/Z sentences may be more difficult to reanalyze because they require syntactic changes that the parser cannot carry out without conscious effort and occasional failure (e.g., Lewis, 1998; Pritchett, 1992), or because it is difficult to identify what changes need to be made in the first place (Van Dyke & Lewis, 2003).

Overall, our results indicate that specialized cognitive mechanisms like triage and reanalysis are needed to fully capture the patterns in human reading and judgment data, in line with the conclusion of Van Schijndel and Linzen (2021). This has implications with regard to a recent trend in psycholinguistics in which researchers attempt to model human language processing using Deep Learning and related computational approaches (Linzen, 2019). Deep Learning models are impressive tools that capture a variety of psycholinguistic effects (Linzen & Baroni, 2021). However, the metric that is most often used to evaluate their performance and to link model predictions to empirical human data is perplexity, which is closely related to surprisal (Goodkind & Bicknell, 2018). The

⁶ Paape and Vasishth (2021) have suggested that some readers may even find comma-less NP/Z sentences impossible to reanalyze, because they think that the comma is obligatory, which would further increase the probability of triage.

fact that surprisal alone cannot account for the magnitude of garden-path effects suggests that Deep Learning models may need to be augmented with additional mechanisms if one wants to use them as realistic models of human sentence processing. One important factor in this regard may be the involvement of conscious awareness in garden-path processing, which has often been invoked in the classical psycholinguistic literature (e.g., Frazier & Rayner, 1982; Marcus, 1980; Pritchett, 1992): In human readers, strong garden paths like the one caused by the NP/Z ambiguity often lead to a conscious experience of processing difficulty, as well as temporary or permanent processing breakdown, and may require deliberate rereading and reanalysis (e.g., Gibson, 1991; Paape & Vasishth, 2022). Language models lack consciousness, which may be part of the reason why they do not fully capture the human data in this specific domain.

Our results also highlight the need for more sophisticated data analysis in psycholinguistics that goes beyond comparing condition means. One fruitful approach that we used in the present work is to establish a linking hypothesis between online measures such as reading times and offline measures such as acceptability judgments or answers to comprehension questions, and to analyze both within a single latent-process model (e.g., Ferreira & Yang, 2019; Nicenboim & Vasishth, 2018; Paape et al., 2021). Another promising avenue is to model reading times and reaction times with mixture distributions (e.g., Lissón et al., 2021; Vasishth, Chopin, Ryder, & Nicenboim, 2017; Vasishth, Jäger, & Nicenboim, 2017), that is, to assume that the data points within conditions may be generated by different processes with different mean latencies and potentially different amounts of variability. By using these approaches, new insights can be gathered from existing data sets. With the public release of research data becoming more and more common, complex computational models can be evaluated with regard to a wealth of different phenomena, and continue to inform theory building (Roberts & Pashler, 2000).

Acknowledgements

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project number 428960187 – PI: Dario Paape.

References

- Altmann, G. T., Garnham, A., & Dennis, Y. (1992). Avoiding the garden path: Eye movements in context. *Journal of Memory and Language*, 31(5), 685–712.
- Bader, M., & Meng, M. (2018). The misinterpretation of noncanonical sentences revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(8), 1286–1311.
- Batchelder, W. H., & Riefer, D. M. (1999). Theoretical and empirical review of multinomial process tree modeling. *Psychonomic Bulletin & Review*, 6(1), 57–86.
- Bixler, R., & D’Mello, S. (2015). Automatic gaze-based detection of mind wandering with metacognitive awareness. In *International Conference on User Modeling, Adaptation, and Personalization* (pp. 31–43).
- Christianson, K., Hollingworth, A., Halliwell, J. F., & Ferreira, F. (2001). Thematic roles assigned along the garden path linger. *Cognitive Psychology*, 42(4), 368–407.
- Clifton, C., Traxler, M. J., Mohamed, M. T., Williams, R. S., Morris, R. K., & Rayner, K. (2003). The use of thematic role information in parsing: Syntactic processing autonomy revisited. *Journal of Memory and Language*, 49(3), 317–334.
- Ferreira, F., & Clifton, C. (1986). The independence of syntactic processing. *Journal of Memory and Language*, 25(3), 348–368.
- Ferreira, F., & Henderson, J. M. (1991). Recovery from misanalyses of garden-path sentences. *Journal of Memory and Language*, 30(6), 725–745.
- Ferreira, F., & Yang, Z. (2019). The problem of comprehension in psycholinguistics. *Discourse Processes*, 56(7), 485–495.
- Fodor, J. D., & Inoue, A. (1994). The diagnosis and cure of garden paths. *Journal of Psycholinguistic Research*, 23(5), 407–434.
- Fodor, J. D., & Inoue, A. (2000). Garden path repair: Diagnosis and triage. *Language and Speech*, 43(3), 261–271.
- 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(5), 992–997.
- Frazier, L. (1979). *On comprehending sentences: Syntactic parsing strategies* (Unpublished doctoral dissertation). University of Connecticut.
- Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, 14(2), 178–210.
- Fujita, H., & Cunnings, I. (2021). Reanalysis processes in non-native sentence comprehension. *Bilingualism: Language and Cognition*, 24(4), 628–641.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis* (Third ed.). Boca Raton, FL: Chapman and Hall/CRC.
- Gibson, E. A. F. (1991). *A computational theory of human linguistic processing: Memory limitations and processing breakdown* (Unpublished doctoral dissertation). Carnegie Mellon University.
- Goodkind, A., & Bicknell, K. (2018). Predictive power of word surprisal for reading times is a linear function of language model quality. In *Proceedings of the 8th Workshop on Cognitive Modeling and Computational Linguistics* (pp. 10–18).
- Grodner, D., Gibson, E., Argaman, V., & Babyonyshev, M. (2003). Against repair-based reanalysis in sentence comprehension. *Journal of Psycholinguistic Research*, 32(2), 141–166.

- Hale, J. (2001). A probabilistic Earley parser as a psycholinguistic model. In *Second Meeting of the North American Chapter of the Association for Computational Linguistics*.
- Heck, D. W., & Erdfelder, E. (2016). Extending multinomial processing tree models to measure the relative speed of cognitive processes. *Psychonomic Bulletin & Review*, 23(5), 1440–1465.
- Hofmeister, P., Jaeger, T. F., Arnon, I., Sag, I. A., & Snider, N. (2013). The source ambiguity problem: Distinguishing the effects of grammar and processing on acceptability judgments. *Language and Cognitive Processes*, 28(1-2), 48–87.
- Jäger, L. A., Engelmann, F., & Vasishth, S. (2017). Similarity-based interference in sentence comprehension: Literature review and Bayesian meta-analysis. *Journal of Memory and Language*, 94, 316–339.
- Klauer, K. C., & Kellen, D. (2018). RT-MPTs: Process models for response-time distributions based on multinomial processing trees with applications to recognition memory. *Journal of Mathematical Psychology*, 82, 111–130.
- Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, 106(3), 1126–1177.
- Lewis, R. L. (1998). Reanalysis and limited repair parsing: Leaping off the garden path. In J. D. Fodor & F. Ferreira (Eds.), *Reanalysis in sentence processing* (pp. 247–285). Dordrecht: Kluwer.
- Linzen, T. (2019). What can linguistics and deep learning contribute to each other? response to pater. *Language*, 95(1), e99–e108.
- Linzen, T., & Baroni, M. (2021). Syntactic structure from deep learning. *Annual Review of Linguistics*, 7, 195–212.
- Lissón, P., Pregla, D., Nicenboim, B., Paape, D., Van het Nederend, M. L., Burchert, F., ... Vasishth, S. (2021). A computational evaluation of two models of retrieval processes in sentence processing in aphasia. *Cognitive Science*, 45(4), e12956.
- Logačev, P., & Vasishth, S. (2016). A multiple-channel model of task-dependent ambiguity resolution in sentence comprehension. *Cognitive Science*, 40(2), 266–298.
- Logačev, P., & Dokudan, N. (2021). A multinomial processing tree model of RC attachment. In *Proceedings of the 11th Workshop on Cognitive Modeling and Computational Linguistics* (pp. 39–47).
- Marcus, M. P. (1980). *A theory of syntactic recognition for natural language*. Cambridge, MA: MIT Press.
- McElree, B. (1993). The locus of lexical preference effects in sentence comprehension: A time-course analysis. *Journal of Memory and Language*, 32(4), 536–571.
- McLachlan, G., & Peel, D. (2000). *Finite mixture models*. New York: Wiley.
- McVay, J. C., & Kane, M. J. (2012). Drifting from slow to “d’oh!”: Working memory capacity and mind wandering predict extreme reaction times and executive control errors. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(3), 525–549.
- Meng, M., & Bader, M. (2000). Mode of disambiguation and garden-path strength: An investigation of subject-object ambiguities in German. *Language and Speech*, 43(1), 43–74.
- Mitchell, D. C., Shen, X., Green, M. J., & Hodgson, T. L. (2008). Accounting for regressive eye-movements in models of sentence processing: A reappraisal of the Selective Reanalysis hypothesis. *Journal of Memory and Language*, 59(3), 266–293.
- Nicenboim, B., Schad, D. J., & Vasishth, S. (2021). *An Introduction to Bayesian Data*

- Analysis for Cognitive Science*.
<https://vasishth.github.io/bayescogsci/book/>.
- Nicenboim, B., & Vasishth, S. (2018). Models of retrieval in sentence comprehension: A computational evaluation using Bayesian hierarchical modeling. *Journal of Memory and Language*, 99, 1–34.
- Ollman, R. (1966). Fast guesses in choice reaction time. *Psychonomic Science*, 6(4), 155–156.
- Paape, D., Avetisyan, S., Lago, S., & Vasishth, S. (2021). Modeling misretrieval and feature substitution in agreement attraction: A computational evaluation. *Cognitive Science*, 45(8), e13019.
- Paape, D., & Vasishth, S. (2021). *When nothing goes right, go left: A large-scale evaluation of bidirectional self-paced reading*. PsyArXiv preprint. Retrieved from psyarxiv.com/d7pvz
- Paape, D., & Vasishth, S. (2022). Is reanalysis selective when regressions are consciously controlled? *Glossa Psycholinguistics*, 1(1).
- Prasad, G., & Linzen, T. (2019). *How much harder are hard garden path sentences than easy ones?* (Poster presented at CogSci 2019)
- Pritchett, B. L. (1992). *Grammatical competence and parsing performance*. Chicago, IL: University of Chicago Press.
- R Core Team. (2021). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <https://www.R-project.org/>
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, 9(3), 438–481.
- Rayner, K., & Sereno, S. C. (1994). Regressive eye movements and sentence parsing: On the use of regression-contingent analyses. *Memory & Cognition*, 22(3), 281–285.
- Reichle, E. D., Reineberg, A. E., & Schooler, J. W. (2010). Eye movements during mindless reading. *Psychological Science*, 21(9), 1300–1310.
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? a comment on theory testing. *Psychological Review*, 107(2), 358–367.
- 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(2), 179–194.
- Smallwood, J. (2011). Mind-wandering while reading: Attentional decoupling, mindless reading and the cascade model of inattention. *Language and Linguistics Compass*, 5(2), 63–77.
- Spivey, M. J., & Tanenhaus, M. K. (1998). Syntactic ambiguity resolution in discourse: modeling the effects of referential context and lexical frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(6), 1521–1543.
- Stahl, C. (2006). Multinomiale Verarbeitungs-Baummodelle in der Sozialpsychologie. *Zeitschrift für Sozialpsychologie*, 37(3), 161–171.
- Stan Development Team. (2020). *RStan: the R interface to Stan*. Retrieved from <http://mc-stan.org/> (R package version 2.21.2)
- Stan Development Team. (2021). Stan modeling language users guide and reference manual [Computer software manual]. Retrieved from <https://mc-stan.org> (Version 2.21.2)
- Sturt, P. (1996). Monotonic syntactic processing: A cross-linguistic study of attachment

- and reanalysis. *Language and Cognitive Processes*, 11(5), 449–494.
- Sturt, P., Pickering, M. J., & Crocker, M. W. (1999). Structural change and reanalysis difficulty in language comprehension. *Journal of Memory and Language*, 40(1), 136–150.
- Swets, B., Desmet, T., Clifton, C., & Ferreira, F. (2008). Underspecification of syntactic ambiguities: Evidence from self-paced reading. *Memory & Cognition*, 36(1), 201–216.
- Tabor, W., & Hutchins, S. (2004). Evidence for self-organized sentence processing: digging-in effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(2), 431–450.
- Trueswell, J. C., Tanenhaus, M. K., & Garnsey, S. M. (1994). Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language*, 33(3), 285–318.
- Ulrich, R. (2009). Uncovering unobservable cognitive mechanisms: The contribution of mathematical models. In F. Rösler, C. Ranganath, B. Röder, & R. H. Kluwe (Eds.), *Neuroimaging of human memory: Linking cognitive processes to neural systems* (pp. 25–41). Oxford: Oxford University Press.
- Unsworth, N., & McMillan, B. D. (2013). Mind wandering and reading comprehension: examining the roles of working memory capacity, interest, motivation, and topic experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(3), 832–842.
- Van Dyke, J. A., & Lewis, R. L. (2003). Distinguishing effects of structure and decay on attachment and repair: A cue-based parsing account of recovery from misanalyzed ambiguities. *Journal of Memory and Language*, 49(3), 285–316.
- Van Gompel, R. P., Pickering, M. J., Pearson, J., & Liversedge, S. P. (2005). Evidence against competition during syntactic ambiguity resolution. *Journal of Memory and Language*, 52(2), 284–307.
- Van Gompel, R. P., Pickering, M. J., & Traxler, M. J. (2000). Unrestricted race: A new model of syntactic ambiguity resolution. In D. H. A. Kennedy R. Radach & J. Pynte (Eds.), *Reading as a perceptual process* (pp. 621–648). Elsevier.
- Van Schijndel, M., & Linzen, T. (2018). Modeling garden path effects without explicit hierarchical syntax. In *Proceedings of CogSci 2018*.
- Van Schijndel, M., & Linzen, T. (2021). Single-stage prediction models do not explain the magnitude of syntactic disambiguation difficulty. *Cognitive Science*, 45(6), e12988.
- Vasishth, S., Chopin, N., Ryder, R., & Nicenboim, B. (2017). Modelling dependency completion in sentence comprehension as a Bayesian hierarchical mixture process: A case study involving Chinese relative clauses. *arXiv preprint arXiv:1702.00564*.
- Vasishth, S., Jäger, L. A., & Nicenboim, B. (2017). Feature overwriting as a finite mixture process: Evidence from comprehension data. *arXiv preprint arXiv:1703.04081*.
- Vasishth, S., Mertzen, D., Jäger, L. A., & Gelman, A. (2018). The statistical significance filter leads to overoptimistic expectations of replicability. *Journal of Memory and Language*, 103, 151–175.
- Vehtari, A., Gabry, J., Yao, Y., & Gelman, A. (2019). *loo: Efficient leave-one-out cross-validation and WAIC for Bayesian models*. Retrieved from <https://CRAN.R-project.org/package=loo> (R package version 2.1.0)
- Vehtari, A., Gelman, A., & Gabry, J. (2015). Pareto smoothed importance sampling. *arXiv preprint arXiv:1507.02646*.
- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using

leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5), 1413–1432.

Zhang, H., Qu, C., Miller, K. F., & Cortina, K. S. (2020). Missing the joke: Reduced rereading of garden-path jokes during mind-wandering. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 46(4), 638–648.