The role of incremental and superficial processing in the depth charge illusion: Experimental and modeling evidence

Dario Paape*

University of Potsdam

Abstract

The depth charge illusion occurs when compositionally incongruous sentences such as *No detail is too unimportant to be left out* are assigned plausible non-compositional meanings (*Don't leave out details*). Results of two online reading and judgment experiments show that moving the incongruous degree phrase to the beginning of the sentence in German (lit. “Too unimportant to be left out is surely no detail”) results in an attenuation of this semantic illusion, implying a role for incremental processing. Two further experiments show that readers cannot consistently turn the communicated meaning of depth charge sentences into its opposite, and that acceptability varies greatly between sentences and subjects, which is consistent with superficial interpretation. A meta-analytic fit of the Wiener diffusion model to data from six experiments shows that world knowledge is a systematic driver of the illusion, leading to stable acceptability judgments. Other variables, such as sentiment polarity, influence subjects’ depth of processing. Overall, the results shed new light on the role of superficial processing on the one hand and of communicative competence on the other hand in creating the depth charge illusion. I conclude that the depth charge illusion combines aspects of being a persistent processing “bug” with aspects of being a beneficial communicative “feature”, making it a fascinating object of study.

Keywords: semantic illusion, superficial processing diffusion model

1 Introduction

Consider the three sentences in (1).

- (1) a. No detail is too unimportant to be left out.
- b. No detail is too unimportant to be left in.
- c. No detail is unimportant enough to be left out.

Sentence (1a) is judged acceptable and sensible by many native speakers of English, and is usually taken to mean *Don't leave out even seemingly unimportant details*. However, given enough time, most speakers

*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.

also agree that (1b) and (1c) are acceptable and sensible. This is surprising, given that *left out* in (1a) is the antonym of *left in* in (1b). Similarly, *too* in (1a) is the logical dual of *enough* in (1c) (*too unimportant* $\approx$ *not important enough*), so that the meaning of the two sentences cannot be the same. The degree phrase *too unimportant to be left out* in (1a) is, in fact, incongruous: Because more important things are usually less likely to be left out, something can be too important to be left out, but nothing can be too *unimportant* to be left out. This semantic incongruity within the *too*-phrase should not be affected by the presence of the negative quantifier *no*: Asserting that no detail has the nonsensical property should not rescue the sentence.¹ In addition, because (1b) means that details should be left in, (1a) should mean that details should be left out, but the salient reading of (1a) is that even seemingly unimportant details should *not* be left out.

In the analysis of Paape et al. (2020), which is based on Meier’s (2003) analysis of the semantics of *too*, (1a) should have the following compositional meaning:

- (2) There is no detail such that the maximal e such that the detail is e -unimportant $>$ the maximal e^* such that, if the detail is e^* -unimportant, it should be left out.

where e is an extent on the “unimportance” scale, and e^* is the threshold on this scale that needs to be exceeded for a detail *not* to be left out. What is crucial here is that the underlying scale is pragmatically odd: One would want a scale on which there is a threshold of “unimportance” beyond which a detail *could* be left out. Furthermore, because the compositional semantics would dictate that the threshold is *not* crossed for any detail, details *should* be left out.

The fact that (1a) receives a compositionally unlicensed interpretation is known as the depth charge illusion, and was first discussed by Wason and Reich (1979). It has since been experimentally investigated across different languages. The illusion occurs in English (O’Connor, 2015; O’Connor, 2017; Zhang et al., 2023), Greek (Natsopoulos, 1985; Giannouli, 2016), Danish (Kizach et al., 2015), and German (Paape et al., 2020). It can also be observed “in the wild”, as shown by corpus data (Cook and Stevenson, 2010; Fortuin, 2014).

The classic account of Wason and Reich (1979) treats the depth charge effect as a processing error: The combination of semantic-pragmatic incongruity and multiple negation in examples like (1a) “overloads” the capacity of the language processing system, which makes readers default to a superficial, plausible interpretation (see also Sanford and Sturt, 2002). The role of negation-induced complexity in the illusion is supported by the fact that removing one or more negative elements from the sentence — the initial *no*, the negative adjective, and/or the negative verb – attenuates the effect (O’Connor, 2015; O’Connor, 2017; Kizach et al., 2015; Paape et al., 2020).

Wason and Reich (1979) did not specify the precise mechanism by which superficial interpretations are computed. One possibility is that readers treat the sentence as an unstructured “bag of words” and use their lexical-associative knowledge as well as general world knowledge to derive a likely meaning (“When these

¹A nonsensical sentence such as *No chair is made of liquid wood* can be said to be true in the sense that the nonsensical property *obviously* does not apply to any chair in existence, but this is not the interpretation that is usually reported for (1a).

words occur together, the speaker will probably intend X”, cf. Paape et al., 2020). The “overloading” account offers an explanation of the depth charge effect that is rooted solely in linguistic performance: The illusion is an aberration due to derailed compositional processing, and does not reflect speakers’ underlying grammatical competence, which should, in principle, enable them to parse the sentence correctly.

Anecdotally, the depth charge effect is remarkably persistent: laypeople and linguistic experts are often unable to see that (1a) is nonsensical when interpreted compositionally, and will sometimes argue vehemently that the initial negation somehow removes the internal incongruity of the degree phrase *too unimportant to be left out*. Many informants also report that even though they understand the problem when it is explained to them, they are still not able to distinguish between novel “bad” and “good” instances of the *No X is too Y to Z* schema.

Given its persistence and subjects’ striking lack of error awareness, is it plausible that the depth charge effect is a performance error? Cook and Stevenson (2010) and Fortuin (2014) have put forward an alternative theory that treats the semantic inversion effect (*left out* → *not left out*) as the skilled use of a linguistic device. The authors argue that *No X is too Y to Z* is a stored grammatical construction or template that can license both the “negative”, inverted meaning in (1a) and the “positive”, uninverted meaning in (1b), depending on context. The *No X is too Y to Z* construction is assumed to reflect a conventionalized, acquired mapping between form and meaning. I will contrast this view with the original *performance error account* and treat it as a subtype of *communicative competence account*. Like the performance error account, this type of account assumes that the depth charge effect cannot be explained compositionally. However, the difference between the two classes of accounts is that the communicative competence view sees the depth charge effect as *purposive* or *rational* in the sense of Chater and Oaksford (1999): The persistent non-compositional interpretation is not seen as an error but as a reflection of linguistic and general communicative skill. I will use the term “communicative competence” here to include any kind of systematic, language-related knowledge beyond grammatical competence in the Chomskyan sense that helps an interlocutor recover a speaker’s or writer’s intended meaning. This includes but is not limited to pragmatic principles, knowledge of idiomatic meanings, and meta-knowledge of how language is processed by an interlocutor (see e.g., Hymes, 1972; Coseriu, 1985; Lehmann, 2007; Rickheit et al., 2008 for discussion).

Under this view, another recent account by Zhang et al. (2023) can also be classed as a communicative competence account, in that the depth charge effect is treated as a rational inference about what the speaker presumably intended to communicate. Instead of treating such pragmatic reasoning as a last-resort rescue mechanism, the account of Zhang et al. treats it as a rational adaptation to the nature of human communication.

Zhang et al. analyze the illusion within the noisy-channel model of Levy (2008) and Gibson et al. (2013). In the noisy-channel model, it is assumed that readers have prior assumptions about sensible sentence meanings, and that they also have a mental model of how errors arise during communication. The processing system combines the prior with the linguistic input and attempts to reconstruct the intended meaning, taking into account the possibility that the input may have been deformed by speech errors, fallible memory, and

other factors. For instance, Zhang et al. (2023) argue that the sentence *No detail is too unimportant to be left out* may be reconstructed as *No detail is so unimportant as to be left out*.

The noisy-channel model is underspecified with regard to whether such “repair” operations rise to consciousness. Given the aforementioned lack of error awareness, it would appear that they do not: Subjects do not give verbal reports to the extent of “I know what this sentence is supposed to mean, but I think there was a mistake, so I mentally corrected it”. The depth charge effect still occurs when participants are explicitly tasked with identifying semantically anomalous utterances and when the sentences are read more than once, which casts some doubt on the role of mental error correction (Paape et al., 2020). On the other hand, the human language processor’s error-correction mechanism may be so sophisticated that it can “repair” certain errors automatically and effortlessly, especially when the error is one that is likely to occur during language production (Frazier and Clifton, 2015).

Even if one wants to maintain processing overload as the driving force behind the depth charge effect, it is highly likely that subjects’ interpretations are at least partly driven by the charitable assumption that most sentences are sensible (e.g., Fillenbaum, 1974). This casts some doubt on the framing of the depth charge effect as an error, in the evaluative sense that the language processing system has “failed” to capture the grammatically “correct” meaning of the sentence. As pointed out by Wason and Reich (1979), it can be argued that “failing” to compute the literal meaning of the utterance is irrelevant as long as the pragmatic meaning is recovered. Taking into account the social dimension of an utterance in addition to its literal content can be seen as a form of rationality (Chase et al., 1998; Hertwig and Gigerenzer, 1999) and of communicative competence. Nevertheless, it is an empirical question whether the depth charge effect in particular is better characterized as a performance-related, potentially maladaptive *bug* in the cognitive system, as originally assumed by Wason and Reich (1979), or rather as a communicative *feature*, either in the form of a grammatical construction or in being the product of a sophisticated error-correction process.

2 Does the illusion depend on a specific linear order?

One interesting question about depth charge sentences regards the interaction between compositionality and incremental processing. In canonical depth charge sentences, the negative quantifier *no* and the degree phrase headed by *too* arrive in a linear order that is at odds with the compositional makeup of the sentence: The adjective and the verb in the degree phrase must be combined into a property (being *too unimportant to be left out*) that is then combined with the initial noun phrase and negated. Such tensions between the order in which the input words are processed and the way they need to be semantically combined are highly common in language processing, and raise important questions about compositionality itself (Baggio et al., 2012; Beck and Tiemann, 2018). A unique aspect of depth charge sentences is that the processor appears to incrementally combine *no* and *too* before the verb of the degree phrase is even encountered. In sentence completion studies, subjects often produce completions such as ... *be left out* or semantically similar continuations for (3a)

(O'Connor, 2015; O'Connor, 2017), but much less often for (3b) (Paape et al., 2020).

- (3) a. No detail is too unimportant to . . . (*be left out, be forgotten*)
- b. Some details are too unimportant to . . . (*be left in, be included*)

Continuations like . . . *be left out* or . . . *be forgotten* are based on a reading of the preamble where *even unimportant details should not be X*, which is not compositionally licensed, as per (2) above. It thus appears that the negative quantifier *no* affects the incremental processing of the degree phrase headed by *too*. In particular, it has been suggested that *too* loses its negative meaning ($X \text{ is too } Y \text{ to } Z \rightarrow X \text{ cannot/should not } Z$) when embedded under *no*, giving it a meaning close to that of *enough* ($X \text{ is } Y \text{ enough to } Z \rightarrow X \text{ can/should } Z$) (O'Connor, 2015; O'Connor, 2017). An alternative account by Fortuin (2014) is that *no* and *too* in combination trigger the appearance of an additional “rhetorical” negation on the verb that emphasizes the construction’s admonitory meaning (*Don’t ignore details!*). Yet another proposal by Paape et al. (2020) is that readers apply a superficial heuristic that combines the negative quantifier *no* and the negative prefix *un-*, so that the negations cancel each other out ($No \text{ detail is too unimportant to be left out} \rightarrow Some/all \text{ details are too important to be left out}$). Importantly, all of the proposed accounts assume that *no* precedes, and presumably also syntactically dominates (c-commands), the degree phrase. But what if the degree phrase is processed *before* the negation is encountered?

Consider the German example in (4), in which the degree phrase has been fronted. Fronting in German is usually analyzed as syntactic movement (e.g. Thiersch, 1978), with the degree phrase’s base position being to the right of *kein Detail*, “no detail”.

- (4) Zu unwichtig, um ausgelassen zu werden, ist sicher kein Detail.
 too unimportant to left.out to get is surely no detail
 ‘Surely, no detail is too unimportant to be left out.’

Does the illusion still arise in this configuration? Fortuin (2014) explicitly argues that *no* needs to be processed before *too* in order for meaning inversion to occur. He assumes that there is a unique, rhetorically licensed interpretation of the substring *No X is too Y* in which *too* loses its negative meaning. This reading then triggers the inverted reading of the entire *No X is too Y to Z* construction with the negated verb. The account of Fortuin (2014) thus predicts that the illusion should be eliminated in (4).

Linear order should also play a role under the overloading account of Wason and Reich (1979): Assuming that readers interpret (4) incrementally, they should have a chance to identify the degree phrase as being incongruous before the global negation is processed and overloads the system. The illusion should thus occur less often in the fronted version, because the superficial processing route is not taken as often as in the canonical version.

The only account that does not straightforwardly predict an effect of linear order is the noisy-channel account of Zhang et al. (2023): Reconstructing the intended meaning of the sentence should not be more difficult with a fronted degree phrase compared to the canonical word order.²

²The experiments reported in this paper were conducted before Zhang et al. (2023) proposed their account, but given the direct

To investigate the effect of linear order inversion on the depth charge illusion, I conducted two experiments that combined online reading measures and acceptability judgments. Assuming that reading times are a signal of processing effort, if superficial processing plays a role in the illusion, one would expect reading to be faster when readers compute superficial interpretations, that is, when they rely on their intuition rather than basing their response on a full compositional parse of the sentence.³

3 Experiments 1 & 2

3.1 Participants

For Experiment 1, 20 native speakers of German were recruited from the student population at the University of Potsdam. For Experiment 2, another 62 native speakers were recruited from the same population. All participants received either 5 € or course credit as compensation.

3.2 Materials

The experimental materials are available at <https://osf.io/2u8p7>. A total of 32 items were constructed according to a 2×2 design with the factors condition (control versus depth charge) and linear order (canonical versus inverted). In the control condition, the negative quantifier *kein*, “no”, on the noun phrase was replaced with *so manche/r/s*, “some (a)”. For Experiment 2, a short preamble as well as a post-critical clause were added to account for processing spillover (Mitchell, 1984). In inverted sentences, an adverb was added after the copula to make the sentences sound more natural.⁴ An example item is shown in Table 1. The experimental sentences were presented using a Latin-square procedure and randomly intermixed with 64 fillers. The fillers were designed to contain a mixture of sensible/correct and incongruous/incorrect sentences. Many fillers featured topicalization constructions similar to the inverted depth charge sentences, and one-fourth of the fillers featured the negative-polarity item *jemals*, “ever”, along with a negative quantifier in either a licensing or a non-licensing position (Drenhaus et al., 2005).

All experimental sentences should be equally incongruous under a compositional reading: the incongruous degree phrase *too unimportant to be left out* appears in all conditions, and the presence or absence of the negative quantifier should not affect its internal incongruity. Higher acceptability ratings in the depth charge condition compared to the control condition would thus indicate that the depth charge effect has occurred, creating an illusion of acceptability.

If readers are more likely to notice the incongruity of the degree phrase in the inverted compared to the canonical conditions, reading times should be higher in the *be left out* region in inverted sentences. On

relevance of the noisy-channel concept to the depth charge phenomenon, I nevertheless include the predictions here. I believe this is unproblematic, as I am not adapting the predictions to the data in any way.

³A feature of the depth charge illusion that *prima facie* appears to be out of line with the performance error account is that cognitive “overload” should slow the reader down rather than speeding them up. However, as discussed by Paape et al. (2020), if the processing breakdown and the subsequent “quick and dirty” interpretation occur at the same point in the sentence and the resulting reading time ends up being faster than under compositional interpretation, an overall speedup is indeed expected.

⁴The results of both experiments indicate that order inversion did not noticeably affect acceptability: The interpolated median rating for canonical sentences was 2.8 in Experiment 1 and 3.2 in Experiment 2, compared to 2.8 and 3.4 for inverted sentences.

Preamble (Experiment 2 only)

Man ◇ bemerkt: ...
one notices

Control, canonical

So manches Detail ◇ ist ◇ zu unwichtig, ◇ um ◇ ausgelassen zu werden,
some a detail is too unimportant to get left.out

Control, inverted

Zu unwichtig, ◇ um ◇ ausgelassen zu werden, ◇ ist ◇ sicher ◇ so manches Detail,
too unimportant to get left.out is surely some a detail

Depth charge, canonical

Kein Detail ◇ ist ◇ zu unwichtig, ◇ um ◇ ausgelassen zu werden,
no detail is too unimportant to get left.out

Depth charge, inverted

Zu unwichtig, ◇ um ◇ ausgelassen zu werden, ◇ ist ◇ sicher ◇ kein Detail,
too unimportant to get left.out is surely no detail

Post-critical clause (Experiment 2 only)

... was ◇ die Künstlerin ◇ besonders ◇ macht.
which the artist special makes

TABLE 1: Example item used in Experiment 1 (whole-sentence reading) and Experiment 2 (self-paced reading). Diamonds (◇) indicate boundaries between presentation regions in Experiment 2.

the other hand, readers have to integrate more material at this point in canonical sentences, so that such the effects may cancel out. Crucially, however, if the initial negation masks the incongruity in the canonical depth charge condition, an interaction should appear in this region, such that canonical but not inverted depth charge sentences show a speedup relative to controls.

3.3 Procedure and data analysis

The experiments were run using Linger (Rohde, 2003). After giving informed consent, participants were instructed to read the sentences at their own pace and indicate after each sentence how acceptable it was on a scale from 1 (not acceptable at all) and 7 (perfectly acceptable). Participants were told to rate a sentence as perfectly acceptable if it was true, coherent and contained no grammatical mistakes. In Experiment 1, the entire sentence was shown on the screen at once. After a key press, the sentence disappeared and participants gave their rating. Experiment 2 used masked self-paced reading (Just et al., 1982). After the final key press, participants gave their rating. Reading times as well as the time taken to assign the rating were recorded.

All data and analysis code, including contrast coding, prior choices and full posterior plots, are available at <https://osf.io/2u8p7>. Data analysis was carried out in R (R Core Team, 2018) using the brms package (Bürkner, 2017), which provides an interface to the Stan language for Bayesian inference (Stan Development Team, 2018). All models were fitted with full variance-covariance matrices (Schielzeth and Forstmeier, 2008; Barr et al., 2013). Ratings were analyzed using cumulative logit models with non-equidistant cutpoints (e.g., Liddell and Kruschke, 2018). For Experiment 1, reading time and rating time were summed to obtain a compound measure of decision time, as the rating process cannot be meaningfully disentangled from the reading process: Given the design of the experiment, it is highly likely that participants start reasoning about

FIGURE 1: Experiment 1. Top: Mean log decision time (reading time + rating time) by condition, with 95% confidence intervals. Bottom: Acceptability ratings by condition, along with interpolated medians.

the acceptability of the sentence as soon as they start reading.

For the analysis of the self-paced reading data, regions were aligned by content rather than linear order. The region following the depth charge clause, which always contained a single word (*was*, “which” in Table 1 above), was analyzed to investigate potential spillover effects. Across all models, I report the posterior mean and 95% credible interval of effects for which more than 95% of the posterior probability are either above or below zero. Zero being included in the confidence interval does not mean that there is evidence for an effect being absent, but only that absence of an effect is plausible given the data and the prior.

3.4 Results

Figure 1 shows decision times (reading time + rating time) and ratings in Experiment 1 by condition. Figure 2 shows reading times, rating times and ratings in Experiment 2 by condition.⁵

3.4.1 Experiment 1 results

In Experiment 1, decision times (reading time plus rating time) were longer for sentences with inverted linear order compared to sentences with canonical linear order ($\hat{\Delta} = 1.8$ s, CrI: [1 s, 2.6 s]). Depth charge sentences were rated more acceptable than control sentences ($\hat{\Delta} = 1.07$, CrI: [0.56, 1.63]). There was also an interaction between condition and linear order ($\hat{\Delta} = -0.75$, CrI: [-1.08, -0.4]), due to depth charge sentences being rated

⁵A plot showing reading times by condition separately for each linear order is available in the online supplementary materials.

FIGURE 2: Experiment 2. Top: Mean log reading times by region and condition, with 95% confidence intervals. Middle: Mean log rating times by region and condition, with 95% confidence intervals. Bottom: Acceptability ratings by condition, along with interpolated medians.

more acceptable than control sentences when linear order was canonical ($\hat{\Delta} = 1.88$, CrI: [1.04, 2.83]) but not convincingly so when linear order was inverted ($\hat{\Delta} = 0.34$, CrI: [-0.17, 0.81]).

The interpolated median rating across all fillers was 3.5. The highest-rated filler sentence (*Some predators are so fast that even birds cannot escape them*) received a median rating of 6.9 while the lowest-rated filler sentence (*Politics is a science that is not practiced by people who only drink juice from bottles*) received a median rating of 1.2.

3.4.2 Experiment 2 results

The reading-time results by region are as follows:⁶

No/some a detail Reading times in this region were longer in inverted sentences compared to canonical sentences ($\hat{\Delta} = 341$ ms, CrI: [234 ms, 453 ms]), presumably because the region appeared later in inverted sentences and more material had to be integrated. Furthermore, reading times were shorter when the noun phrase contained the negative quantifier *no* as opposed to *some (a)* ($\hat{\Delta} = -123$ ms, CrI: [-190 ms, -55 ms]), possibly because an additional scalar implicature was computed for *some*. (e.g., Tomlinson et al., 2013).

too unimportant Reading times were shorter in inverted sentences compared to canonical sentences ($\hat{\Delta} = -55$ ms, CrI: [-113 ms, 3 ms]), presumably because less material had to be integrated at this point in inverted sentences, and longer in depth charge versus control sentences ($\hat{\Delta} = 41$ ms, CrI: [1 ms, 80 ms]).⁷

be left out Reading times were shorter in inverted sentences compared to canonical sentences ($\hat{\Delta} = -76$ ms, CrI: [-161 ms, 15 ms]), again presumably due to differences in integration costs.

[spillover] At the first word after the end of the depth charge clause, reading times were shorter in depth charge sentences compared to control sentences ($\hat{\Delta} = -30$ ms, CrI: [-56 ms, -2 ms]). There was also an interaction ($\hat{\Delta} = 51$ ms, CrI: [20 ms, 82 ms]), due to faster reading times in depth charge sentences compared to control sentences when linear order was canonical ($\hat{\Delta} = -81$ ms, CrI: [-130 ms, -32 ms]) but an opposite tendency when linear order was inverted ($\hat{\Delta} = 22$ ms, CrI: [-8 ms, 51 ms]).

Rating times were longer for depth charge sentences than for control sentences ($\hat{\Delta} = 108$ ms, CrI: [-11 ms, 224 ms]). As in Experiment 1, depth charge sentences were rated more acceptable than control sentences ($\hat{\Delta} = 1.23$, CrI: [0.63, 1.8]), and there was an interaction between condition and linear order ($\hat{\Delta} = -0.61$, CrI: [-0.92, -0.32]), due to depth charge sentences being rated more acceptable than control sentences when linear order was canonical ($\hat{\Delta} = 1.84$, CrI: [1.17, 2.49]) but less clearly so when linear order was inverted ($\hat{\Delta} = 0.56$, CrI: [-0.04, 1.2]).

The interpolated median rating across all fillers was 3.5.

3.5 Discussion

The results of Experiments 1 and 2 suggest that the depth charge effect is reduced when the incongruous degree phrase linearly precedes the negation. When linear order was canonical, depth charge sentences were rated more acceptable than control sentences, despite both sentence types being compositionally incongruous. By contrast, when linear order was inverted through fronting of the degree phrase, the difference in ratings was much smaller.

The results are inconclusive with regard to whether the effect of word order on the illusion is qualitative or quantitative. Complete absence of the illusion in inverted sentences would be consistent with the view

⁶There was no indication of effects in the two one-word regions, *ist*, “is”, and *zu*, “too”.

⁷The latter finding is surprising given that *too unimportant* was encountered before *no/some* in the inverted conditions, but as Figure 2 shows, the effect is mainly driven by the canonical conditions.

that the negation needs to be processed before the degree phrase in order for meaning inversion to occur, as proposed by Fortuin (2014). Irrespective of whether the difference between the word orders is qualitative or quantitative, the fact that there *is* a difference casts doubt on the error-correction account of Zhang et al. (2023), under which linear order should not affect interpretation.⁸ There was, however, a numerical tendency for depth charge sentences to be rated more acceptable than control sentences even with inverted linear order. In Experiment 2, about one fourth of inverted depth charge sentences were rated 6 or higher, despite the incongruous degree phrase appearing sentence-initially. The existence of high ratings in the inverted condition is not consistent with the strong claim by Fortuin (2014) that the *No X is too Y to Z* construction can only receive its “negative” meaning when the negation precedes the degree phrase. High ratings in the presence of an (arguably) obvious inconsistency right at the beginning of the sentence are more in line with previous findings showing that readers often do not evaluate sentences incrementally on a word-by-word basis but use a more global “gist” strategy that ignores local inconsistencies (e.g., Barton and Sanford, 1993; Kamas et al., 1996; Hannon and Daneman, 2001), and are thus more consistent with a performance error account as opposed to a communicative competence account.

In the reading time analyses, a speedup was observed for canonical depth charge sentences in the spillover region in Experiment 2. Earlier experimental findings on semantic anomaly detection suggest that processing is faster when the anomaly is missed compared to when it is noticed (Bohan and Sanford, 2008; Cook et al., 2018). The observed speedup could thus index superficial, non-compositional processing. Assuming that superficial processing is also reflected in higher ratings, there may be a connection between the two dependent measures. However, adding reading times at the spillover region as a predictor of the end-of-sentence acceptability rating yielded no strong indication of an effect ($\hat{\Delta} = 0.14$, CrI: $[-0.52, 0.87]$); there is thus no indication that faster reading in this region coincides with a stronger depth charge illusion. Nevertheless, it is striking that sentences with an additional negation are processed faster than their non-negated counterparts at the point where the entire *No X is too Y to Z* construction has been processed, and that the effect is limited to the canonical depth charge configuration.

Overall, the experimental results show that the linear order of the degree phrase and the negative quantifier is an important factor in the depth charge illusion, which is unexpected under the error-correction account of Zhang et al. (2023) but expected under Wason and Reich’s (1979) “overloading” account. A reduction of the depth charge effect in inverted sentences is also broadly consistent with Fortuin’s (2014) claim that the canonical linear order triggers “rhetorical” negation in depth charge sentences. However, the fact that there were occasional high acceptability ratings for inverted depth charge sentences is inconsistent with a strong interpretation of Fortuin’s claim that “[the compositional interpretation] cannot be undone by expressing [negation] after it has already received this interpretation” (p. 278).

⁸A reviewer suggests that a plausible production error may be the “blending” of two constructions, namely *No detail should be left out* and *No detail is too unimportant to be included*. Such “blendings” have been discussed by Frazier and Clifton (2015), who argue that the sentence processor may be able to automatically “repair” them. This is an intriguing possibility, but as for the account of Zhang et al. (2023), it is unclear why linear order would affect the processor’s repair abilities.

4 Is the interpretation underspecified?

Given that the incongruity of the degree phrase is sometimes ignored even when it appears sentence-initially, what kind of internal representation do participants generate in such trials? Under the view that *No X is too Y to Z* is a stored construction with a specific meaning (Fortuin, 2014), once a “negative” instance of the construction has been identified, participants should settle on the meaning and mentally store the appropriate proposition. Presumably, this should also occur under the error-correction account of Zhang et al. (2023) once the reader has reconstructed the intended message. Under both accounts, readers should thus come to believe that the depth charge sentence expresses a well-defined proposition (e.g., *Details, no matter how seemingly unimportant they may be, should not be left out*).

It is less clear how things play out under the “overloading” view. If the processor attempts compositional processing, fails, and finally resorts to superficial heuristics to derive a plausible meaning, are such “good enough” representations distinguishable from compositionally-derived representations? For instance, it might be that readers get only a vague impression of what the sentence means (*Something about not leaving out details*), or that well-specified meaning is stored along with a mental “flag” that encodes some lingering doubt about its correctness. On the other hand, even if processing is superficial, the resulting representations may nevertheless be detailed and readers may be strongly committed to their interpretation. It is thus important to investigate if there is a correspondence between readers’ impression of having grasped the sentence’s meaning and the objective quality of their internal representation.

When the depth charge illusion occurs, it is often accompanied by a strong “feeling of knowing”: Speakers are usually convinced that they are interpreting the sentence correctly. Kizach et al. (2015) had their Danish-speaking participants indicate their subjective confidence in their interpretation of depth charge sentences (“Was your answer a guess?”), finding that subjects responded “no” in about 80% of cases. However, feelings of knowing have been shown to be negatively correlated with accuracy for stimuli whose “consensual” answer – that is, the answer given by the majority of subjects – is wrong (Koriat, 1975, 2008). Thus, if superficial processing is the root cause of the depth charge illusion and readers *could*, in principle, make an effort to process the sentence more deeply, they would have no incentive to do so, given their faith in the illusory interpretation.

In the sentence-processing literature, some authors have argued that rather than being a last resort, as assumed by the “overloading” theory of Wason and Reich (1979), superficial processing is ubiquitous, because it saves the reader time and effort (Ferreira and Patson, 2007). Some accounts assume that “surface” or “gist” semantics are always computed before any detailed syntactic structure is built (e.g., Townsend and Bever, 2001), while other accounts assume that surface semantics and compositional, syntax-mediated semantics are computed in parallel (e.g., Kim and Osterhout, 2005; Kuperberg, 2007). As an example, for simple, canonical sentences such as *The dog bit the man*, the assignment of thematic roles is so plausible and straightforward that no detailed semantic evaluation may be necessary. In line with this view, non-canonical sentences like *The dog was bitten by the man* are often misinterpreted (Ferreira, 2003), presumably because

the language processor sticks to its ingrained interpretation habits and compositional processing is aborted once the current representation is deemed “good enough” by some standard. In such cases, surface semantics will determine the representation of the sentence (e.g., Karimi and Ferreira, 2016).

The idea of “good enough” processing is strongly related to dual-process theories of reasoning (Ferreira et al., 2009; Christianson, 2016), and to Simon’s concept of “satisficing” (Simon, 1955, 1956), which assumes that decisions are optimized for maximum payoff, but only up to an aspiration level. If accurate, detailed comprehension takes too much time, the aspiration level is lowered and processing becomes “good enough”.⁹ It is thus implied that “good enough” interpretations are somehow impoverished compared to full compositional interpretations. Nevertheless, previous research suggests that participants are convinced that their “good enough” interpretations are sufficient for answering comprehension questions with relatively high confidence, even if the answers are often incorrect (Christianson et al., 2001).

“Good enough” processing is also strongly related to the concept of underspecification. Underspecification of meaning occurs when certain aspects of a semantic representation are simply not computed during processing, or are left in a “fuzzy” state in which multiple meanings are possible (Pinkal, 1996). There is some evidence that downstream processing becomes more effortful with such underspecified meanings. For instance, in a self-paced reading study on globally ambiguous versus disambiguated sentences (*The son/maid of the princess who scratched herself in public . . .*), Swets et al. (2008) found that detailed end-of-trial comprehension questions (*Did the maid scratch in public?*) were more difficult to answer for globally ambiguous sentences, presumably because participants did not settle on an interpretation prior to the appearance of the question.¹⁰ Similar findings are reported by Dwivedi (2013) for sentences with ambiguous quantifier scope. There is also evidence that processing is slowed when underspecified or comparatively less specified linguistic representations are accessed by downstream retrieval triggers such as verbs or syntactic gaps (Hofmeister, 2011; Hofmeister and Vasishth, 2014; Paape et al., 2018).¹¹

If depth charge sentences like *No detail is too unimportant to be left out* are processed superficially, their representations may be underspecified, in the sense that it is left undecided whether unimportant details should or should not be left out, and perhaps even if the sentence makes any sense at all. Subjecting these representations to downstream processing operations should then lead to difficulty, especially if these operations require fully specified propositions as input.

One test case is to try and mentally invert the meaning of a depth charge sentence (“What if I want to say the opposite of *No detail is too unimportant to be left out*?”). If depth charge sentences have an underspecified semantics, it should be difficult to do this, because it may not be entirely clear to the reader what the intended proposition is. “Difficult” here means that in an experimental setting, readers should be

⁹Bever and Townsend (2001) refer to such superficial representations as “drafts” that may or may not be replaced by “real”, deep comprehension if compositional processing is allowed to finish. Similarly, Karimi and Ferreira (2016) refer to “interim outputs”, which may be “refined if necessary” (p. 1019).

¹⁰See Logačev and Vasishth (2016) for a critical re-evaluation of this claim, and for an additional distinction between partial and complete underspecification.

¹¹A lot of research has also been done on the semantic underspecification of homonymous and polysemous words, see Frisson (2009) for a review. I am, however, not aware of any work in this context where the underspecified meaning is later accessed by an additional task or linguistic trigger.

slower to mentally invert depth charge sentences than compositionally well-formed control sentences, and there should be more trials in which they have to resort to guessing whether the inverted meaning of a depth charge sentence is sensible or not. By contrast, if an abstract *No X is too Y to Z* construction exists, it should be relatively straightforward to compute and store the intended meaning, and also to invert it: If *No detail is too unimportant to be left out* encodes the meaning *Don't leave out details*, turning this meaning into its opposite should yield *Leave out details*.

The construction-based account also plausibly predicts a relative *stability* of meaning within individual sentences and within participants: At least for some sentences, the intended rhetorical meaning should be highly salient for all participants because of shared world knowledge (Fortuin, 2014). Furthermore, assuming that participants apply their acquired linguistic knowledge of the construction across different experimental trials, meaning should also be relatively stable within participants. By contrast, the “good enough” approach predicts relatively high inter-individual variability for depth charge sentences, compared to compositionally sensible sentences: In different subsets of trials, the parser may compute the incongruous compositional meaning, the superficial, sensible meaning, or no well-defined meaning at all. This should result in a more noisy and thus more heterogeneous interpretation profile. Additionally, underspecification and unsystematic processing breakdowns should lead to an increased probability of guessing, which should also contribute to more heterogeneous interpretations.

I have intentionally avoided the term “negation”, as I want to establish a novel experimental task that does not rely on linguistic negation. Given that depth charge sentences already contain several negative elements, it would be too demanding for participants to add yet another one. Instead, in Experiments 3 and 4, participants were asked to mentally invert the meaning of depth charge and control sentences. Specifically, participants were told to interpret the sentences in a world where people always say the opposite of what they mean. The manipulation thus does not target the actual *utterance* in question but rather the reader’s mental representation of the utterance’s meaning, the presumably intended *message* (Hagoort et al., 1999), which may or may not be clear. In contrast to Experiments 1 and 2, Experiments 3 and 4 used compositionally sensible control sentences to check if the manipulation was successful: Inverting the meaning of a compositionally sensible sentence such as *Some details are too important to be left out* should presumably be easy for participants and thus yield a reasonable baseline, but a compositionally incongruous control sentence from Experiments 1 and 2 such as *Some details are too unimportant to be left out* would likely result in confusion. By contrast, if depth charge sentences receive a stable and well-specified illusory interpretation, they should behave similarly to compositionally sensible control sentences in that they should be easy to “invert”.

5 Experiments 3 & 4

5.1 Participants

For Experiment 3, 20 native speakers of German were recruited from the student population at the University of Potsdam. For Experiment 4, another 89 native speakers were recruited from the same population. All participants received course credit as compensation.

5.2 Materials

The materials are available at <https://osf.io/2u8p7>. The same sentences used in Experiments 1 and 2 were adapted to a new 2×2 design with the factors condition (control versus depth charge) and interpretation world (normal world versus NEG-world). Control sentences in Experiments 3 and 4 were compositionally sensible, which was achieved by changing the adjective in the degree phrase to its antonym. Table 2 shows an example item. As before, 64 fillers were used.

Control

Manch ein	Detail	ist	zu wichtig,	um	ausgelassen zu werden.
some a	detail	is	too important	to	get left.out

Depth charge

Kein	Detail	ist	zu unwichtig,	um	ausgelassen zu werden.
no	detail	is	too unimportant	to	get left.out

TABLE 2: Example item used in Experiments 3 and 4.

5.3 Procedure and data analysis

Experiment 3 was run in the laboratory using Linger while Experiment 4 was run online on the Ibex farm (Drummond, 2018). The procedure in Experiment 3 was mostly analogous to that used in Experiment 1, but acceptability was judged in a binary fashion (*Is the sentence sensible, correct and true?*) rather than on a graded scale. After the judgment, a second prompt appeared randomly after 50% of trials, asking participants if their judgment had been a guess. In Experiment 4, judgments were recorded on the same screen on which the sentence was presented as opposed to a separate screen, and guessing was not probed. At the beginning of each experiment, participants were instructed that half of the judgments would be made in the NEG-world, where people always say the opposite of what they mean. As an example, participants were told that the sentence *I like Brussels sprouts* would mean *I don't like Brussels sprouts* in the NEG-world.¹² Participants were also told that some sentences may not be sensible or true in either world. During the experiment, trials belonging to the different worlds were visually distinguished: For the normal world, sentences and judgment prompts were presented in black font on a white background, and the note “Normal world” appeared at the

¹²Linguistic negation was used in the example to illustrate the task, not to instruct participants to mentally add a negation to every sentence. However, a subset of participants nevertheless reported using this strategy (see below).

FIGURE 3: Experiment 3. Top left: Mean decision times (reading time + rating time) by condition, with 95% confidence intervals. Top right: Percentage of “acceptable” judgments by condition. Bottom: Percentage of self-reported guesses by condition.

top of the screen. For the NEG-world, sentences and judgment prompts were presented in white font on a black background, and the note “NEG-world” appeared at the top of the screen.¹³

5.4 Results

Figure 3 shows decision times (reading time + rating time) and the percentage of positive acceptability judgments and self-reported guesses in Experiment 3 by condition. Figure 4 shows decision times and the percentage of positive acceptability judgments in Experiment 4 by condition.

5.4.1 Experiment 3 results

Acceptability judgments were made more slowly in the NEG-world compared to the normal world ($\hat{\Delta} = 4$ s, CrI: [2.6 s, 5.5 s]). Acceptability judgments were also made more slowly for depth charge sentences compared

¹³Incidentally, the visual cues resembled those used in Carlson’s (1989) study, which compared linguistic and nonlinguistic negation of logic gates. The validity of the current design is supported by the fact that the latency patterns observed by Carlson were similar to the ones in Experiments 3 and 4.

FIGURE 4: Experiment 4. Left: Mean decision times (reading time + rating time) by condition, with 95% confidence intervals. Right: Percentage of “acceptable” judgments by condition.

to control sentences ($\hat{\Delta} = 2.3$ s, CrI: [1.3 s, 3.2 s]), and more quickly for “accept” than for “reject” judgments ($\hat{\Delta} = -0.7$ s, CrI: [-1.5 s, 0.1 s]). There was also a three-way interaction between world, condition, and judgment ($\hat{\Delta} = -1$ s, CrI: [-2 s, 0.1 s]): For “accept” judgments only, the latency difference between worlds was larger for control sentences ($\hat{\Delta} = 5.9$ s, CrI: [4 s, 7.8 s]) than for depth charge sentences ($\hat{\Delta} = 2.5$ s, CrI: [0.15 s, 4.9 s]).

Fewer positive acceptability judgments were given in the NEG-world compared to the normal world ($\hat{\Delta} = -40\%$, CrI: [-61%, -14%]). Fewer positive acceptability judgments were also given in the depth charge condition compared to the control condition ($\hat{\Delta} = -23\%$, CrI: [-39%, -7%]). There was also an interaction ($\hat{\Delta} = 41\%$, CrI: [25%, 55%]): Control sentences received fewer positive acceptability judgments in the NEG-world compared to the normal world ($\hat{\Delta} = -67\%$, CrI: [-83%, -43%]) but there was no indication of a difference for depth charge sentences ($\hat{\Delta} = 1\%$, CrI: [-26%, 27%]).

Guessing percentages showed no reliable indication of a difference between worlds, conditions and answer types, though there was a numerical tendency towards a three-way interaction, with fewer guesses for positive judgments of control sentences in the normal world compared to all other judgments (see Figure 3). This tendency can be mapped onto the decision time results: “Accept” judgments for control sentences in the normal world were made especially quickly and with especially high certainty, compared to all other condition/judgment pairs. Judgment certainty may thus be partly reflected by judgment latency.

5.4.2 Experiment 4 results

The results largely aligned with those of Experiment 3. Acceptability judgments were made more slowly in the NEG-world compared to the normal world ($\hat{\Delta} = 2.3$ s, CrI: [1.6 s, 2.9 s]). Acceptability judgments were also made more slowly for depth charge sentences compared to control sentences ($\hat{\Delta} = 1.4$ s, CrI: [0.9 s, 1.9 s]), and more quickly for “accept” than for “reject” judgments ($\hat{\Delta} = -0.5$ s, CrI: [-1 s, -0.1 s]). There was also a three-way interaction between world, condition, and judgment ($\hat{\Delta} = -0.7$ s, CrI: [-1.2 s, -0.2 s]): For “accept” judgments only, the latency difference between worlds was larger for control sentences ($\hat{\Delta} = 4.1$ s, CrI: [3 s, 5.2 s]) than for depth charge sentences ($\hat{\Delta} = 1.5$ s, CrI: [0.3 s, 2.7 s]). As in Experiment 3, this interaction presumably reflects the high certainty of “accept” judgments for control sentences in the normal world.

Fewer positive acceptability judgments were given in the NEG-world compared to the normal world ($\hat{\Delta} = -33\%$, CrI: [-43%, -22%]). Fewer positive acceptability judgments were also given in the depth charge condition compared to the control condition ($\hat{\Delta} = -25\%$, CrI: [-33%, -17%]). There was also an interaction ($\hat{\Delta} = 43\%$, CrI: [35%, 51%]): Control sentences received fewer positive acceptability judgments in the NEG-world compared to the normal world ($\hat{\Delta} = -66\%$, CrI: [-75%, -56%]), while depth charge sentences showed some indication of the opposite tendency ($\hat{\Delta} = 12\%$, CrI: [-3%, 26%]).

5.5 Discussion

The results of Experiments 3 and 4 show that participants largely succeeded at inverting the meaning of compositionally sensible control sentences (*Some details are too important to be left out*) in the NEG-world. In Experiment 4, which had the larger sample size, control sentences were judged acceptable in about 86% of trials in the normal world, compared to about 32% of trials in the NEG-world. By contrast, in the depth charge conditions (*No detail is too unimportant to be left out*), there were more positive judgments in the NEG-world (47%) than in the normal world (38%). The fact that the difference between worlds was much smaller for depth charge compared to control sentences is consistent with superficial processing of depth charge sentences, resulting in underspecified meanings.

By contrast, the pattern in decision times is, at face value, not consistent with underspecification of meaning in the depth charge condition: If underspecified meanings are more difficult to invert, the latency difference between the normal world and the NEG-world should have been larger for depth charge sentences than for control sentences, but the opposite pattern was observed. One way to reconcile the latency pattern with earlier findings on underspecification is to assume that there is a maximum amount of time that participants are willing to spend on a given sentence (Paape et al., 2020). Indeed, the 95% credible interval for the mean acceptability of depth charge sentences in the NEG-world includes 50% (CrI: [38%, 54%]), which is consistent with participants losing interest and giving a random judgment. Yet, in Experiment 3, subjects’ self-reported guessing percentage for depth charge sentences was at around 22% in the normal world, consistent with the estimate reported by Kizach et al. (2015), and around 29% in the NEG-world. While an increased guessing

rate in the inverted world may have pushed the overall acceptability of depth charge sentences closer to 50% in Experiment 4, the majority of judgments were confident, indicating that the data do not reflect pure guessing.

The high proportion of confident judgments is somewhat unexpected given that incorrect intuitive responses have been found to coincide with reduced subjective confidence (De Neys et al., 2011). This correlation suggests that participants usually have some degree of insight into the shaky basis of their judgments (De Neys and Bonnefon, 2013). One reason for the high levels of subjective confidence in depth charge sentences could be the high accessibility of indirect cues to meaning, such as world knowledge (*Details are often important*), which may lead to an overconfidence effect (Koriat, 2012). When participants activate additional knowledge that is unrelated to a given task, this can result in increased judgment confidence while lowering objective accuracy (Halberstadt and Levine, 1999; Hall et al., 2007). Participants may be under the impression that they have given a well-justified judgment, because they think they know what the sentences are *supposed* to mean, while ignoring their compositional structure.

5.6 Is there evidence for consistent meaning inversion?

The experimental task used in Experiment 4 was presumably challenging for participants, which may have led to variation in the judgments. Indeed, seven participants reported during debriefing that they had been unsure as to which parts of the sentence they were supposed to invert or “negate” in the NEG-world. It is thus informative to see to what extent different participants performed the task as intended, and to what extent participants succeeded at mentally inverting the meaning of different sentences in the NEG-world. One intriguing possibility is that the 50/50 pattern observed for depth charge sentences across worlds is the result of stability within sentences but variation across sentences: Half of the experimental sentences could be accepted in 100% of cases in the normal world but rejected in 100% of cases in the NEG-world, while the other half could show the opposite pattern.

A stimulus sentence that has a stable illusory interpretation and that is perfectly “invertible” would show 100% positive judgments in the normal world and 100% negative judgments in the NEG-world for both the control and depth charge conditions: For instance, the sentence *No detail is too unimportant to be left out* could mean *Don’t leave out details* in the normal world and be acceptable for all participants, but would mean *Leave out details* in the NEG-world and be unacceptable for all participants. Figure 5 graphs the percentage of positive acceptability judgments in the normal world versus the percentage of negative acceptability judgments in the NEG-world across the control and depth charge conditions for each participant and each sentence.¹⁴

As the plots show, no sentence in the sample is perfectly “invertible”. Especially at the participant level, the data are relatively noisy, and neither depth charge sentences nor control sentences line up on the diagonal. However, control sentences show less variability than depth charge sentences, which are distributed across the entire area of the plot. This could mean that subjects were often confused and resorted to guessing in the depth charge condition, especially in the NEG-world. What does *not* seem to be the case is that participants

¹⁴Due to the Latin-square design, each participant encountered each sentence in only one condition, so that sentences cannot be compared across conditions within individuals and vice versa.

FIGURE 5: Experiment 4. Percentage of positive acceptability judgments in the normal world plotted against the percentage of negative acceptability judgments in the NEG-world, by condition. Top: Percentages by item, based on data from 89 subjects. Bottom: Percentages by subject, based on data from 32 items.

reliably identified a stored grammatical construction and were then able to invert its meaning (*No detail is too unimportant to be left out* → *Don't leave out details* → *Leave out details*): If this was the case, the sentence means should have landed on the diagonal in Figure 5, because the acceptance rate of a particular depth charge sentence in the normal world should directly translate into its rejection rate in the NEG-world.¹⁵

If subjects resorted to guessing in the depth charge condition, the variation across different sentences and participants should be entirely random. However, as the rate of self-reported guessing in Experiment 3 was relatively low, this seems unlikely. There may thus be systematic factors driving the variability. The variability between subjects cannot be further investigated based on the current data, given that no individual-differences measures were collected, but such distinguishing measures can be gathered a posteriori from corpora for the experimental sentences. The final section of the paper investigates a number of such measures in the context of a process model whose parameters can be linked to “good enough” processing on the one hand and communicative competence on the other.

¹⁵Proponents of the construction-based approach may argue that participants had trouble identifying the intended meaning of the ambiguous *No X is too Y to Z* construction, given that no context was provided, and that this uncertainty was the source of the variability in judgments. In future work, a manipulation of the discourse context could help shed light on this question.

6 Modeling the influence of sentence-level cues on the depth charge effect

For the present investigation, I selected three factors that may contribute to the depth charge effect: World knowledge, superficial semantic cohesion, and sentiment polarity. World knowledge has been investigated as a factor in the depth charge illusion by several authors (see Paape et al., 2020 for a review). Sentiment polarity has also been argued to be a factor in the interpretation, given that depth charge sentences tend to contain many words associated with negative sentiments, such as *unimportant* and *leave out* (e.g., Cook and Stevenson, 2010; Paape et al., 2020). Superficial semantic cohesion refers to the idea that lexical associations between the content words may lead to faster formation of a “gist” meaning, which may in turn interfere with analytical processing (Kuperberg, 2007).

Combining the data from all four experiments reported above, as well as the publicly available data from Experiment 1 (whole-sentence reading) and Experiment 2A (eye tracking) of Paape et al. (2020), I thus conducted an exploratory meta-analytic investigation of item-level cues to meaning in depth charge sentences. The resulting data set contained data from a total of 270 subjects.

The combined data were fitted with the Wiener diffusion model for two-choice reaction times (e.g., Ratcliff, 1978; Ratcliff and Smith, 2004). The modeling code is available at <https://osf.io/2u8p7>. In the Wiener diffusion model, the decision process – that is, to accept or to reject the sentence – is a noisy process of evidence accumulation that can terminate at either of two decision boundaries. The two advantages of the diffusion model are that its parameters can be interpreted directly, and that the decision outcomes and their accompanying latencies are modeled in conjunction rather than separately. Out of the four parameters in the brms implementation of the model (Bürkner, 2019; see also Wabersich and Vandekerckhove, 2014), two are of interest here:

- **Drift rate:** The rate at which the process of evidence accumulation approaches one of the boundaries. Drift can be positive (towards acceptance) or negative (towards rejection).
- **Boundary separation:** The distance between the decision boundaries. When the boundaries are further apart, more evidence needs to be accumulated before a choice is made, and premature responses are less likely. This parameter can thus be interpreted as reflecting response caution and/or response confidence, with higher separation corresponding to slower but more cautious, confident responses (e.g., Zeguers et al., 2011).

Figure 6 shows a schematic illustration of the diffusion process. As the plot shows, when the decision boundaries are close together (dashed lines), noise will occasionally result in a boundary being crossed, leading to a premature decision. When the boundaries are further apart (solid lines), premature decisions are less likely.

For the present purpose, evidence accumulation is equated with reading (and potentially rereading) the

FIGURE 6: Illustration of the diffusion process. The boundary separation parameter determines the distance between “accept” and “reject” boundaries.

sentence and reaching a boundary is equated with pressing the response button to give the accept/reject judgment. Across all experiments, the cumulative time elapsed between initial sentence presentation and giving the rating or judgment was used as the reaction time for the trial.

Negative effects on the boundary separation parameter would indicate more shallow processing: When the decision boundaries are closer together, minor fluctuations in the process of evidence accumulation lead to fast but premature decisions. Reduced boundary separation has been linked to reduced cognitive monitoring (Dutilh et al., 2012; Huff and Aschenbrenner, 2018) and to impulsive responding (Hedge et al., 2020). Boundary separation decreases after engaging in high-effort tasks, suggesting that subjects set a lower criterion when their cognitive resources are depleted (Lin et al., 2020). Given these associations, it is plausible to link the boundary separation parameter, which represents subjects’ aspiration level or depth of processing, to “good enough” interpretation: When the boundaries are close together, reading will often be aborted prematurely and the judgment will be based on incomplete evidence.

By contrast, positive changes to the drift rate – that is, faster drift towards acceptance – would indicate that participants make strategic use of non-compositional cues, and that these cues are given the status of evidence for an intended interpretation. One such cue is world knowledge about plausible sentence meanings. If prior plausibility affects the speed of evidence accumulation in favor of accepting the depth charge sentence, judgments should become faster and be less influenced by noise when world knowledge is strong. At the level of representation, such strategic processing would presumably result in stable meaning representations and would be in line with communicative competence accounts like those proposed by Cook and Stevenson (2010), Fortuin (2014), and Zhang et al. (2023).

Importantly, it should be noted that depending on one’s conceptualization of “good enough” processing, effects of world knowledge on the drift parameter may be compatible with the framework: If it is assumed that superficial interpretation heuristics represent a systematic adaptation to the pressures of everyday language

processing and to imperfections in the input (Ferreira and Patson, 2007; Christianson, 2016), the “good enough” framework turns into another variant of the communicative competence account, under which readers rationally combine different sources of evidence to infer meaning. However, in contrast with the rational inference account of Zhang et al. (2023), the “good enough” framework emphasizes that the use of superficial heuristics is the result of limited temporal, motivational and attentional resources. This implies that there is some notion of an “optimal” processing route that *would* be taken if these resources were unlimited. It is thus a *lack* of “optimal” processing and the use of “satisficing” strategies (Ferreira et al., 2009) that distinguishes the “good enough” account from communicative competence accounts, which is reflected in the predicted effects on the boundary separation parameter that represents the reader’s aspiration level.

The four item-level predictors were operationalized as follows:

World knowledge Paape et al. (2020) conducted an ancillary study in which participants were asked how strongly they agreed with compositionally sensible versions of the sentences, e.g., *Some details are too important to be left out*, on a scale from 1 (disagree completely) to 5 (agree completely). The mean agreement value for each item reflects the strength of world knowledge about the illusory meaning of the relevant depth charge sentence (*No detail is too unimportant to be left out* → *Don’t leave out details*).

Superficial semantic cohesion Semantic cohesion was operationalized as the semantic similarity between the noun, the adjective and the verb, specifically the similarity of their vector-based meaning representations as computed from large corpora. Two measures were computed for each sentence: A local measure based on word embeddings from a 5-gram model (Word2Vec; Mikolov et al., 2013a), and a global measure based on document-level latent semantic analysis (LSA; Martin and Berry, 2007). For both models, the underlying assumption is that words that occur in similar contexts have similar meanings (Landauer and Dumais, 1997). For the 5-gram model, pre-trained embeddings based on the German Wikipedia were used (Yamada et al., 2020). For the LSA model, the pre-trained semantic space of Günther et al. (2015) based on the web-crawled deWaC corpus (Baroni et al., 2009) was used. For both models, the 300-dimensional representations of the noun and the adjective were combined through element-wise addition, under the assumption that this operation reflects semantic integration (Mikolov et al., 2013b). For each experimental item, cosine similarity was then computed between the context (*detail + unimportant*) and the verb (*left out*) as a measure of cohesion. The view of cohesion taken here is broad and not limited to “classical” lexical relations (Morris and Hirst, 2004), and the computation is intentionally limited to content words, under the assumption that readers may extract a “topical gist” from the sentences as a basis for meaning recovery (Landauer, 2007).

Sentiment polarity The sentiment polarity of each experimental sentence was computed using the model of Guhr et al. (2020), which was trained on social media posts and reviews that human annotators had preclassified as “positive”, “negative”, or “neutral”. The model is based on Wolf et al.’s (2019) implementation of BERT (Devlin et al., 2018). For each depth charge sentence, only the content words (noun, adjective, verb) were

considered. The intuition behind the use of sentiment polarity is that depth charge sentences usually invoke a negative presupposition (*One may think that details can be small enough to be left out*), which is negated (*... but one should not leave them out anyway*; see also Cook and Stevenson, 2010; Fortuin, 2014). Participants may thus use their knowledge about the sentiment associated with the relevant lexical items as a heuristic to infer the overall meaning of the sentence. Instead of classifying the sentences as having “positive”, “negative” or “neutral” sentiment, the log odds of the “negative” category were extracted for each item to obtain a continuous predictor. Out of 32 experimental sentences, 26 were classified as negative by the model, while the remaining six were classified as neutral.

Slopes for the four item-level predictors were added to the drift rate and boundary separation parameters of the diffusion model. The predictor values for each sentence are available at <https://osf.io/2u8p7>.¹⁶

The diffusion model was specified hierarchically. Subjects were assumed to have different drift rates, biases, non-decision times, and boundary separations. The bias parameter determines whether the evidence accumulation process is shifted towards one of the boundaries even before the stimulus is encountered. For instance, some subjects may have a general preference to accept or reject sentences. The non-decision time parameter accounts for differences in reaction times that are unrelated to the stimulus, such as slower reading speed or slower key presses. The bias and drift parameters were also assumed to vary between experiments, as the mode of presentation as well as the type and overall amount of materials presented differed between studies. A slope for the interpretation world manipulation was added to the bias parameter, as the instruction to mentally invert the meaning could bias participants towards either acceptance or rejection.

In order to make the studies comparable, only the depth charge condition was modeled. The factors linear order (canonical versus inverted) and interpretation world (normal versus NEG-world) were added to the data sets for all experiments. The experiments of Paape et al. (2020) contained neither manipulation, and thus only contributed data to the [canonical order, normal world] condition. Graded judgments were transformed to binary judgments by considering only the extreme ends of the 1–7 scale: Ratings of 1 or 2 were treated as rejections, ratings of 6 or 7 were treated as endorsements, and the remaining ratings were excluded from the analysis.¹⁷ In order to keep the number of parameters manageable and to facilitate interpretation, interaction terms were only added for the drift rate parameter, and only two-way interactions were included. Effects in the form of two-way interactions on the boundary separation parameter are possible in principle, but would be difficult to interpret theoretically: For instance, it is hypothetically possible that response caution is less affected by word order when world knowledge is strong, but such an effect is, to my knowledge, not straightforwardly predicted by any existing processing theory. This is because response caution is a factor that has not received much attention in psycholinguistics (but see Hammerly et al., 2019). Nevertheless, it should be kept in mind that constraining the model in the aforementioned way also constrains the interpretation of

¹⁶Data for one sentence were removed prior to analysis because one of the critical words was not present in the vocabulary of the LSA corpus. For a small subset of the remaining sentences, some words were replaced with close neighbors (e.g., *injury* instead of *head injury*) for the same reason.

¹⁷This, of course, results in “middling” ratings being ignored in the analysis, thus limiting its coverage of the data. In future work, it would be worthwhile to repeat Experiments 1 and 2 with binary judgments, in order to allow for a direct mapping of the responses.

FIGURE 7: Coefficient estimates from the diffusion model analysis. sep = boundary separation; drift = drift rate; senti = sentiment polarity; LSA = latent semantic analysis (global cohesion); W2V = Word2Vec (local cohesion); WK = world knowledge. Boundary separation is on the log scale, bias is on the log-odds scale. Slopes correspond to a unit increase in standard deviation from the mean.

the results.

6.1 Results

Figure 7 shows the coefficient estimates from the diffusion model analysis.

The results can be summarized as follows: Boundary separation — that is, response caution — increases in the NEG-world, when the degree phrase is fronted, and when global semantic cohesion is high. By contrast, strong negative sentiment polarity and strong world knowledge decrease response caution. The drift rate parameter is most strongly affected by linear order and world knowledge: Inverted linear order (*Too unimportant to be left out is surely no detail*) pushes the accumulation process towards rejection while strong world knowledge pushes it toward acceptance. There is also an interaction between world knowledge and interpretation world: Table 3 shows that strong world knowledge results in higher acceptance rates in the normal world but lower acceptance rates in the NEG-world.

World	World knowledge	p(accept)
normal	strong	0.61
normal	weak	0.41
NEG	strong	0.42
NEG	weak	0.50

TABLE 3: Proportions of positive judgments by interpretation world and world knowledge strength.

6.2 Discussion

The diffusion model fit suggests that both factors related to communicative competence and factors related to “good enough” processing influence the interpretation of depth charge sentences. The strongest predictors

of drift rate, and therefore the speed and direction of evidence accumulation, were linear order and world knowledge. During processing, the appearance of the incongruous degree phrase (*too unimportant to be left out*) at the beginning of the sentence provides evidence that the depth charge sentence should be rejected, compared to the canonical version in which the degree phrase appears after the negation. By contrast, strong world knowledge (*Details should not be left out*) is used by participants as evidence in favor of accepting the sentence as sensible.

When the incongruous degree phrase appears sentence-initially, readers do not only accumulate more evidence in favor of rejection but also make more careful judgments, as evidenced by increased boundary separation. This suggests a role for superficial processing in the canonical *No X is too Y to Z* configuration. At the same time, the effect on the drift rate parameter is partly compatible with the claim by Fortuin (2014) that the canonical ordering of the negation and the degree phrase is a necessary condition for the illusion, even though some depth charge sentences are consistently accepted in the inverted condition as well.¹⁸

The effect of world knowledge depends on the interpretation world, suggesting that the underlying mechanism is quite sophisticated: In the normal world, participants tended to endorse depth charge sentences whose illusory interpretation they agreed more with, while in the NEG-world, they tended to endorse the sentences whose illusory interpretation they agreed *less* with.

It is not surprising that readers construct plausible meanings by tapping into their knowledge about the real world, given that world knowledge is also habitually recruited in compositionally well-formed sentences (e.g., Cook and Guéraud, 2005; Cook and O'Brien, 2014; Isberner and Richter, 2013; von der Malsburg et al., 2020). The systematic contribution of world knowledge to meaning implies that the use of this particular source of information is a kind of communicative competence (Coseriu, 1985; Lehmann, 2007; Fortuin, 2014). It is also broadly in line with the error-correction account of Zhang et al. (2023), which assumes that readers attempt to reconstruct the intended meaning of the depth charge sentence, though it remains unclear why the reconstruction process would be affected by linear order.

Based on earlier work, one could have hypothesized that global semantic cohesion should speed up the decision process, given its connection to contextual constraint (Pynte et al., 2008b,a). However, the cohesion measure used here was intended to be superficial and thus considered only content words, ignoring the negative quantifier and the presence of the degree word *too*. Speculatively, participants may have been more careful in judging high-cohesion sentences because the maximally superficial reading (*detail + unimportant + leave out*) was at odds with the illusion reading, which requires the verb to be negated (*Don't leave out even seemingly unimportant details*).

In summary, the modeling results suggest that the depth charge effect may be driven partly by “good enough” processing and partly by communicative competence. The effect of world knowledge shows that the presumably intended meaning of an utterance is sometimes more important to readers than its actual compositional makeup. This is in line with accounts that highlight the social-communicative function of

¹⁸Please refer to the supplementary materials for a plot showing the empirical acceptance rates across conditions by item, as well as the predictions of the diffusion model.

the depth charge construction (Cook and Stevenson, 2010; Fortuin, 2014; Zhang et al., 2023). At the same time, the results suggest that superficial, “good enough” processing is involved in the illusion, as evidenced by the observed effects of world knowledge and sentiment polarity on participants’ aspiration levels. This is in line with the “good enough”, performance-based view, which highlights the role of effort-saving, heuristic processing mechanisms (Wason and Reich, 1979; Paape et al., 2020).

One obvious limitation of the present exploratory investigation is that it was limited to the depth charge condition. Further research using a variety of matched control sentences is needed to find out whether the observed patterns are unique to the illusion configuration or whether the same mechanisms also apply to other types of sentences.

7 General discussion

The experimental part of the present work focused on two issues relating to the depth charge illusion that had not previously been investigated: The role of the linear ordering of the negative quantifier and the degree phrase headed by *too*, and the question of whether readers’ interpretations of depth charge sentences are demonstrably underspecified or “good enough” (e.g., Karimi and Ferreira, 2016). The computational modeling part investigated the influence of sentence-specific, non-compositional cues such as sentiment polarity and superficial semantic cohesion. The overall aim was to compare accounts that analyze the depth charge effect as being due to a failure of compositional processing (Wason and Reich, 1979; Paape et al., 2020) with accounts that highlight the role of readers’ communicative competence, either in the form of knowledge of a particular grammatical construction (Cook and Stevenson, 2010; Fortuin, 2014) or of knowledge about plausible sentence meanings and possible transmission errors (Zhang et al., 2023).

7.1 The effect of linear order, and implications for the competing theories

In Experiments 1 and 2, the depth charge effect in acceptability ratings was reduced when the canonical order *No X is too Y to Z* was changed to *Too Y to Z is surely no X* via syntactic fronting in German. This effect implies that there is a processing component to the depth charge illusion: The internal incongruity of the degree phrase is masked to a larger extent if the negation is processed before the degree phrase. The results are inconclusive with regard to whether the average effect in the non-canonical configuration is zero or just numerically smaller than in the canonical configuration. If the meaning of the depth charge construction crucially depends on the linear ordering of *too* and the negation that scopes over it Fortuin (2014), a purely quantitative effect would be surprising, as the illusion should be completely eliminated in the non-canonical configuration. The effect of linear order is also surprising under the error-correction account of Zhang et al. (2023), as subjects’ ability to reconstruct the intended meaning of the sentence should presumably be unaffected by the rearrangement of the lexical items.

A quantitative effect of linear precedence falls out more naturally under the performance error view, which

includes the “overloading” approach of Wason and Reich (1979): In the canonical case, the incremental interpretation of the negative quantifier followed by *too* triggers a failure of compositional processing, which results in the adoption of a superficial interpretation. In the inverted sentences, this type of overload should be less likely, unless interpretation is delayed until after the negation. This may occur stochastically in some trials but not in others, which would explain why inverted sentences sometimes receive high ratings.

7.2 The effect of the interpretation world, the relationship with underspecification, and the role of sentence-level cues

The goal of Experiments 3 and 4 was to investigate if readers generate underspecified, “shallow” or “good enough” interpretations of depth charge sentences. There is no consensus in the existing literature as to how the shallowness of a representation can be assessed. One widely-used approach uses incorrect answers to comprehension questions as evidence that processing must have been “good enough” (e.g., Christianson et al., 2001; Ferreira, 2003). This approach can be challenged on the grounds that answering comprehension questions may involve additional, error-prone processes beyond those involved in online sentence interpretation (Bader and Meng, 2018; Qian et al., 2018), and that discarded misanalyses may leave memory traces that can influence offline interpretation (Slattery et al., 2013). Due to these limitations, the current experiments tested the shallowness of the representation indirectly: Experiments 3 and 4 required readers to interpret depth charge sentences in the “NEG-world”, where all statements mean the opposite. The results suggest that semantic “inversion” was possible and mostly systematic for sensible control sentences but not for matched depth charge sentences, for which judgments varied widely across subjects and items. The variability in the judgments can be taken as support for “good enough” processing, and thus the performance error view of the illusion.

Another way to distinguish superficial and “deep” aspects of semantic processing is to relate reaction times and judgments in an implemented cognitive model such as the Wiener diffusion model. An exploratory investigation using data from six experiments shed additional light on the roles of “good enough” processing and communicative competence in the depth charge illusion. Depth charge sentences with stronger associated world knowledge showed more consistent “invertibility” across interpretation worlds, as evidenced by the effect of word knowledge on the rate of evidence accumulation. The fact that world knowledge also interacted with the interpretation world suggests that readers used this cue very systematically. The use of world knowledge to inform semantic processing can thus be seen as a type of communicative competence exhibited by the interpretation system. By contrast, superficial semantic cohesion and sentiment polarity both influenced readers’ aspiration levels: Strong negative sentiment polarity, that is, the presence of strongly “negative” words such as *unimportant*, reduced the aspiration level, indicating reduced response caution and more superficial processing, while global superficial cohesion increased the aspiration level, suggesting more cautious processing. Reduced aspiration levels are more in line with the performance error account of the depth charge illusion, given that processing of the stimulus will often remain incomplete.

The modeling results also show that moving the incongruous degree phrase to the beginning of the sentence pushes the accumulation process towards rejection and increases participants' aspiration level, leading to more careful responses. The increased tendency to reject the sentences is broadly compatible with Fortuin's (2014) argument that the depth charge construction can only receive its "rhetorical" interpretation if *no* appears before *too*, though the data show that the constraint is by no means absolute. The effect of linear order on the aspiration level is more straightforwardly accounted for by assuming that the incongruity of the degree phrase is more salient when it is fronted, and is thus more in line with the performance error view as opposed to the communicative competence view.

It is clear that speakers often ignore semantic inconsistency when interpreting and judging depth charge sentences, a tendency that can be reduced by making the inconsistency more salient. There is also a lot of variability between sentences and speakers in terms of whether the illusion occurs, which suggests that the incongruity is noticed in many trials. Furthermore, the fact that speakers mostly fail to mentally invert the meaning of depth charge sentences, or judge the sentences to be equally acceptable whether the meaning is inverted or not, shows that the meaning of *No X is too Y to Z* sentences is often not clear. On the other hand, the pattern of illusory interpretations is far from random: The guiding effect of world knowledge, in particular, points towards a potentially highly useful and "rational" mechanism that brings the reader's semantic expectations to bear on the interpretation process (Gibson et al., 2013; Zhang et al., 2023). The resulting gist can be seen as "more advanced than relying on rote representations of reality" (Reyna, 2021, p. 2): Being able to extract intended meanings beyond the verbatim content of an utterance is presumably a highly important skill in everyday communication. Overall, the answer to the question of whether the depth charge effect is more of a cognitive "bug" in the form of a performance error or more of a cognitive "feature" driven by readers' communicative competence appears to be: a combination of both.

Acknowledgments

All experiments were funded by the University of Potsdam. The author would like to thank Shravan Vasishth, Titus von der Malsburg, the Vasishth Lab team, and the audience at AMLaP 2020 for helpful comments and suggestions. Thanks also go to Johanna Thieke for assistance with data collection.

Conflict of interest

The author has no conflict of interest to declare.

References

- Bader, M. and Meng, M. (2018). The misinterpretation of noncanonical sentences revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(8):1286–1311.
- Baggio, G., Van Lambalgen, M., and Hagoort, P. (2012). The processing consequences of compositionality. In Hinzen, W., Machery, E., and Werning, M., editors, *The Oxford Handbook of Compositionality*, pages 655–672. Oxford University Press.
- Baroni, M., Bernardini, S., Ferraresi, A., and Zanchetta, E. (2009). The wacky wide web: A collection of very large linguistically processed web-crawled corpora. *Language Resources and Evaluation*, 43(3):209–226.
- Barr, D. J., Levy, R., Scheepers, C., and Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3):255–278.
- Barton, S. B. and Sanford, A. J. (1993). A case study of anomaly detection: Shallow semantic processing and cohesion establishment. *Memory & Cognition*, 21(4):477–487.
- Beck, S. and Tiemann, S. (2018). Towards a model of incremental composition. In *Proceedings of Sinn und Bedeutung*, volume 21, pages 143–162.
- Bever, T. G. and Townsend, D. J. (2001). Some sentences on our consciousness of sentences. In Dupoux, E., editor, *Language, Brain, and Cognitive Development: Essays in Honor of Jacques Mehler*, pages 143–155. MIT Press, Cambridge, MA.
- Bohan, J. and Sanford, A. (2008). Semantic anomalies at the borderline of consciousness: An eye-tracking investigation. *Quarterly Journal of Experimental Psychology*, 61(2):232–239.
- Bürkner, P.-C. (2019). Bayesian item response modeling in R with brms and Stan. *arXiv preprint arXiv:1905.09501*.
- Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80(1):1–28.
- Carlson, R. A. (1989). Processing nonlinguistic negation. *The American Journal of Psychology*, 102(2):211–224.
- Chase, V. M., Hertwig, R., and Gigerenzer, G. (1998). Visions of rationality. *Trends in Cognitive Sciences*, 2(6):206–214.
- Chater, N. and Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, 3(2):57–65.
- Christianson, K. (2016). When language comprehension goes wrong for the right reasons: Good-enough, underspecified, or shallow language processing. *Quarterly Journal of Experimental Psychology*, 69(5):817–828.
- Christianson, K., Hollingworth, A., Halliwell, J. F., and Ferreira, F. (2001). Thematic roles assigned along the garden path linger. *Cognitive Psychology*, 42(4):368–407.
- Cook, A. E. and Guéraud, S. (2005). What have we been missing? The role of general world knowledge in discourse processing. *Discourse Processes*, 39(2–3):265–278.

- Cook, A. E. and O'Brien, E. J. (2014). Knowledge activation, integration, and validation during narrative text comprehension. *Discourse Processes*, 51(1-2):26–49.
- Cook, A. E., Walsh, E. K., Bills, M. A., Kircher, J. C., and O'Brien, E. J. (2018). Validation of semantic illusions independent of anomaly detection: Evidence from eye movements. *Quarterly Journal of Experimental Psychology*, 71(1):113–121.
- Cook, P. and Stevenson, S. (2010). No sentence is too confusing to ignore. In *Proceedings of the 2010 Workshop on NLP and Linguistics: Finding the Common Ground*, pages 61–69, Uppsala, Sweden. Association for Computational Linguistics.
- Coseriu, E. (1985). Linguistic competence: What is it really? *The Modern Language Review*, 80(4):XXV–XXXV.
- De Neys, W. and Bonnefon, J.-F. (2013). The ‘whys’ and ‘whens’ of individual differences in thinking biases. *Trends in Cognitive Sciences*, 17(4):172–178.
- De Neys, W., Cromheeke, S., and Osman, M. (2011). Biased but in doubt: Conflict and decision confidence. *PloS one*, 6(1):e15954.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Drenhaus, H., Saddy, D., and Frisch, S. (2005). Processing negative polarity items: When negation comes through the backdoor. In Kepser, S. and Reis, M., editors, *Linguistic evidence: Empirical, theoretical, and computational perspectives*, pages 145–165.
- Drummond, A. (2018). *Ibex farm*. <http://spellout.net/ibexfarm/>.
- Dutilh, G., van Ravenzwaaij, D., Nieuwenhuis, S., van der Maas, H. L., Forstmann, B. U., and Wagenmakers, E.-J. (2012). How to measure post-error slowing: A confound and a simple solution. *Journal of Mathematical Psychology*, 56(3):208–216.
- Dwivedi, V. D. (2013). Interpreting quantifier scope ambiguity: Evidence of heuristic first, algorithmic second processing. *PloS one*, 8(11):e81461.
- Ferreira, F. (2003). The misinterpretation of noncanonical sentences. *Cognitive Psychology*, 47(2):164–203.
- Ferreira, F., Engelhardt, P. E., and Jones, M. W. (2009). Good enough language processing: A satisficing approach. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society*.
- Ferreira, F. and Patson, N. D. (2007). The ‘good enough’ approach to language comprehension. *Language and Linguistics Compass*, 1(1–2):71–83.
- Fillenbaum, S. (1974). Pragmatic normalization: Further results for some conjunctive and disjunctive sentences. *Journal of Experimental Psychology*, 102(4):574–578.
- Fortuin, E. (2014). Deconstructing a verbal illusion: The ‘No X is too Y to Z’ construction and the rhetoric of negation. *Cognitive Linguistics*, 25(2):249–292.
- Frazier, L. and Clifton, C. (2015). Without his shirt off he saved the child from almost drowning: interpreting an uncertain input. *Language, Cognition and Neuroscience*, 30(6):635–647.

- Frisson, S. (2009). Semantic underspecification in language processing. *Language and Linguistics Compass*, 3(1):111–127.
- Giannouli, V. (2016). A verbal illusion reexamined. *Acta Neuropsychologica*, 14(4):323–329.
- Gibson, E., Bergen, L., and Piantadosi, S. T. (2013). Rational integration of noisy evidence and prior semantic expectations in sentence interpretation. *Proceedings of the National Academy of Sciences*, 110(20):8051–8056.
- Guhr, O., Schumann, A.-K., Bahrman, F., and Böhme, H. J. (2020). Training a broad-coverage German sentiment classification model for dialog systems. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 1627–1632.
- Günther, F., Dudschig, C., and Kaup, B. (2015). LSAfun – An R package for computations based on Latent Semantic Analysis. *Behavior Research Methods*, 47(4):930–944.
- Hagoort, P., Brown, C. M., and Osterhout, L. (1999). The neurocognition of syntactic processing. pages 273–316. Oxford University Press, New York.
- Halberstadt, J. B. and Levine, G. M. (1999). Effects of reasons analysis on the accuracy of predicting basketball games. *Journal of Applied Social Psychology*, 29(3):517–530.
- Hall, C. C., Ariss, L., and Todorov, A. (2007). The illusion of knowledge: When more information reduces accuracy and increases confidence. *Organizational Behavior and Human Decision Processes*, 103(2):277–290.
- Hammerly, C., Staub, A., and Dillon, B. (2019). The grammaticality asymmetry in agreement attraction reflects response bias: Experimental and modeling evidence. *Cognitive Psychology*, 110:70–104.
- Hannon, B. and Daneman, M. (2001). Susceptibility to semantic illusions: An individual-differences perspective. *Memory & Cognition*, 29(3):449–461.
- Hedge, C., Powell, G., Bompas, A., and Sumner, P. (2020). Self-reported impulsivity does not predict response caution. *Personality and Individual Differences*, 167:110257.
- Hertwig, R. and Gigerenzer, G. (1999). The ‘conjunction fallacy’ revisited: How intelligent inferences look like reasoning errors. *Journal of Behavioral Decision Making*, 12(4):275–305.
- Hofmeister, P. (2011). Representational complexity and memory retrieval in language comprehension. *Language and Cognitive Processes*, 26(3):376–405.
- Hofmeister, P. and Vasishth, S. (2014). Distinctiveness and encoding effects in online sentence comprehension. *Frontiers in Psychology*, 5:1237.
- Huff, M. J. and Aschenbrenner, A. J. (2018). Item-specific processing reduces false recognition in older and younger adults: Separating encoding and retrieval using signal detection and the diffusion model. *Memory & Cognition*, 46(8):1287–1301.
- Hymes, D. H. (1972). On communicative competence. In Pride, J. and Holmes, J., editors, *Sociolinguistics: Selected Readings*, pages 269–293. Penguin, Harmondsworth.
- Isberner, M.-B. and Richter, T. (2013). Can readers ignore implausibility? Evidence for nonstrategic

- monitoring of event-based plausibility in language comprehension. *Acta Psychologica*, 142(1):15–22.
- Just, M. A., Carpenter, P. A., and Woolley, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General*, 111(2):228–238.
- Kamas, E. N., Reder, I. M., and Ayers, M. S. (1996). Partial matching in the Moses illusion: Response bias not sensitivity. *Memory & Cognition*, 24(6):687–699.
- Karimi, H. and Ferreira, F. (2016). Good-enough linguistic representations and online cognitive equilibrium in language processing. *The Quarterly Journal of Experimental Psychology*, 69(5):1013–1040.
- Kim, A. and Osterhout, L. (2005). The independence of combinatory semantic processing: Evidence from event-related potentials. *Journal of Memory and Language*, 52(2):205–225.
- Kizach, J., Christensen, K. R., and Weed, E. (2015). A verbal illusion: Now in three languages. *Journal of Psycholinguistic Research*, 45:1–16.
- Koriat, A. (1975). Phonetic symbolism and feeling of knowing. *Memory & Cognition*, 3(5):545–548.
- Koriat, A. (2008). Subjective confidence in one’s answers: the consensuality principle. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(4):945–959.
- Koriat, A. (2012). The self-consistency model of subjective confidence. *Psychological Review*, 119(1):80–113.
- Kuperberg, G. R. (2007). Neural mechanisms of language comprehension: Challenges to syntax. *Brain Research*, 1146:23–49.
- Landauer, T. K. (2007). LSA as a theory of meaning. In AT. Landauer, D. McNamara, D. S. and Kintsch, W., editors, *Handbook of Latent Semantic Analysis*, pages 3–34. Erlbaum, Mahwa, NJ.
- Landauer, T. K. and Dumais, S. T. (1997). A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2):211.
- Lehmann, C. (2007). Linguistic competence: Theory and empiry. *Folia Linguistica*, 41(3–4):223–278.
- 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*, pages 234–243.
- Liddell, T. M. and Kruschke, J. K. (2018). Analyzing ordinal data with metric models: What could possibly go wrong? *Journal of Experimental Social Psychology*, 79:328–348.
- Lin, H., Saunders, B., Friese, M., Evans, N. J., and Inzlicht, M. (2020). Strong effort manipulations reduce response caution: A preregistered reinvention of the ego-depletion paradigm. *Psychological Science*, 31(5):531–547.
- Logačev, P. and Vasishth, S. (2016). Understanding underspecification: A comparison of two computational implementations. *Quarterly Journal of Experimental Psychology*, 69(5):996–1012.
- Martin, D. I. and Berry, M. W. (2007). Mathematical foundations behind Latent Semantic Analysis. In T. K. Landauer, D. S. McNamara, S. D. and Kintsch, W., editors, *Handbook of Latent Semantic Analysis*, pages 35–56. Lawrence Erlbaum Associates, Mahwah, NJ.

- Meier, C. (2003). The meaning of *too*, *enough*, and *so... that*. *Natural Language Semantics*, 11(1):69–107.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013a). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013b). Distributed representations of words and phrases and their compositionality. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems 26 (NIPS 2013)*, pages 3111–3119.
- Mitchell, D. C. (1984). An evaluation of subject-paced reading tasks and other methods for investigating immediate processes in reading. In Kieras, D. E. and Just, M. A., editors, *New methods in reading comprehension research*, pages 69–89. Erlbaum, Hillsdale, NJ.
- Morris, J. and Hirst, G. (2004). Non-classical lexical semantic relations. In *Proceedings of the Computational Lexical Semantics Workshop at HLT-NAACL 2004*, pages 46–51.
- Natsopoulos, D. (1985). A verbal illusion in two languages. *Journal of Psycholinguistic Research*, 14(4):385–397.
- O’Connor, E. (2017). The accidental ambiguity of inversion illusions. In Lamont, A. and Tetzloff, K., editors, *Proceedings of NELS 47*, pages 329–342.
- O’Connor, E. (2015). *Comparative illusions at the syntax-semantics interface*. PhD thesis, University of Southern California.
- Paape, D., Hemforth, B., and Vasishth, S. (2018). Processing of ellipsis with garden-path antecedents in French and German: Evidence from eye tracking. *PloS one*, 13(6):e0198620.
- Paape, D., Vasishth, S., and von der Malsburg, T. (2020). Quadruplex Negatio Invertit? The On-Line Processing of Depth Charge Sentences. *Journal of Semantics*, 37(4):509–555.
- Pinkal, M. (1996). Vagueness, ambiguity, and underspecification. In *Semantics and Linguistic Theory*, volume 6, pages 185–201.
- Pynte, J., New, B., and Kennedy, A. (2008a). A multiple regression analysis of syntactic and semantic influences in reading normal text. *Journal of Eye Movement Research*, 2(1):1–11.
- Pynte, J., New, B., and Kennedy, A. (2008b). On-line contextual influences during reading normal text: A multiple-regression analysis. *Vision Research*, 48(21):2172–2183.
- Qian, Z., Garnsey, S., and Christianson, K. (2018). A comparison of online and offline measures of good-enough processing in garden-path sentences. *Language, Cognition and Neuroscience*, 33(2):227–254.
- R Core Team (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Version 4.0.3.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2):59–108.
- Ratcliff, R. and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111(2):333–367.
- Reyna, V. F. (2021). A scientific theory of gist communication and misinformation resistance, with implications

- for health, education, and policy. *Proceedings of the National Academy of Sciences*, 118(15):e1912441117.
- Rickheit, G., Strohner, H., and Vorweg, C. (2008). The concept of communicative competence. *Handbook of communication competence*, 1:15–62.
- Rohde, D. (2003). Linger. <http://tedlab.mit.edu/~dr/Linger/>.
- Sanford, A. J. and Sturt, P. (2002). Depth of processing in language comprehension: Not noticing the evidence. *Trends in Cognitive Sciences*, 6(9):382–386.
- Schielzeth, H. and Forstmeier, W. (2008). Conclusions beyond support: Overconfident estimates in mixed models. *Behavioral Ecology*, 20(2):416–420.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2):129–138.
- Slattery, T. J., Sturt, P., Christianson, K., Yoshida, M., and Ferreira, F. (2013). Lingering misinterpretations of garden path sentences arise from competing syntactic representations. *Journal of Memory and Language*, 69(2):104–120.
- Stan Development Team (2018). *Stan Modeling Language Users Guide and Reference Manual, Version 2.21.0*.
- Swets, B., Desmet, T., Clifton, C., and Ferreira, F. (2008). Underspecification of syntactic ambiguities: Evidence from self-paced reading. *Memory & Cognition*, 36(1):201–216.
- Thiersch, C. L. (1978). *Topics in German syntax*. PhD thesis, Massachusetts Institute of Technology.
- Tomlinson, J. M., Bailey, T. M., and Bott, L. (2013). Possibly all of that and then some: Scalar implicatures are understood in two steps. *Journal of Memory and Language*, 69(1):18–35.
- Townsend, D. J. and Bever, T. G. (2001). *Sentence comprehension: The integration of habits and rules*. MIT Press, Cambridge, MA.
- von der Malsburg, T., Poppels, T., and Levy, R. P. (2020). Implicit gender bias in linguistic descriptions for expected events: The cases of the 2016 United States and 2017 United Kingdom elections. *Psychological Science*, 31(2):115–128.
- Wabersich, D. and Vandekerckhove, J. (2014). The RWiener package: An R package providing distribution functions for the Wiener diffusion model. *The R Journal*, 6(1):49–56.
- Wason, P. C. and Reich, S. S. (1979). A verbal illusion. *The Quarterly Journal of Experimental Psychology*, 31(4):591–597.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., et al. (2019). Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Yamada, I., Asai, A., Sakuma, J., Shindo, H., Takeda, H., Takefuji, Y., and Matsumoto, Y. (2020). Wikipedia2Vec: An efficient toolkit for learning and visualizing the embeddings of words and entities

- from Wikipedia. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 23–30. Association for Computational Linguistics.
- Zeguers, M. H., Snellings, P., Tijms, J., Weeda, W. D., Tamboer, P., Bexkens, A., and Huizenga, H. M. (2011). Specifying theories of developmental dyslexia: A diffusion model analysis of word recognition. *Developmental Science*, 14(6):1340–1354.
- Zhang, Y., Ryskin, R., and Gibson, E. (2023). A noisy-channel approach to depth-charge illusions. *Cognition*, 232:105346.