

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

What processing instructions do connectives provide? Modeling the facilitative effect of the connective

Permalink

<https://escholarship.org/uc/item/2sc1k7pf>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Marchal, Marian

Scholman, Merel Cleo Johanna

Sanders, Ted

et al.

Publication Date

2024

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

What processing instructions do connectives provide?

Modeling the facilitative effect of the connective

Marian Marchal (marchal@coli.uni-saarland.de)

Saarland University, Saarbrücken, Germany

Merel Scholman (m.c.j.scholman@coli.uni-saarland.de)

Saarland University, Saarbrücken, Germany

Utrecht University, Utrecht, the Netherlands

Ted Sanders (t.j.m.sanders@uu.nl)

Utrecht University, Utrecht, the Netherlands

Vera Demberg (vera@coli.uni-saarland.de)

Saarland University, Saarbrücken, Germany

Abstract

Connectives like ‘because’ are referred to as ‘processing instructions’ as they facilitate processing of linguistic material directly following the connective. In an expectation-driven account of discourse processing, this can be attributed to predictions that readers make about the upcoming discourse *relation*, but also to predictions about upcoming discourse *content*. By modeling these two accounts, termed the *relation prediction account* and the *content prediction account* respectively, we show that they make different predictions about when the presence of a connective is most beneficial. In a self-paced reading study, we replicate the facilitative effect of the connective on processing, but do not find any evidence that this effect can be explained by a strong or weak version of either of the two accounts. This suggests that the role of the connective goes above and beyond informing the reader about the upcoming relation and content and possibly triggers a different processing strategy.

Keywords: connectives; coherence; information value; self-paced reading; surprisal; predictability

Introduction

During reading and listening, people create rich mental representations from the linguistic information. A crucial element of such coherent mental representations are *discourse relations*: semantic-pragmatic links between clauses and sentences, such as CAUSE and CONTRAST. Connectives and linking phrases, like ‘because’ and ‘nevertheless’, play an important role in this process. They are often referred to as ‘processing instructions’, informing the reader how to integrate the upcoming clause with the preceding one (Britton, 1994; Gernsbacher, 1997). Connectives facilitate on-line processing during reading (Cozijn, Noordman, & Vonk, 2011; Sanders & Noordman, 2000; van Silfhout, Evers-Vermeul, & Sanders, 2015): The region directly following the connective, below in bold, is read faster in the presence of a connective, (1a), compared to when no connective is present, (1b).

- (1) a. Ana was tired. Even so, she **went to a party** with Jia.
- b. Ana was tired. She **went to a party** with Jia.

This initial speed-up is attributed to the connective informing the reader how the clauses relate to each other, also

referred to as “propositional integration”. In this paper, we investigate two expectation-driven accounts that could explain this facilitative effect of the connective, termed the *relation prediction account* and the *content prediction account*.

The facilitated integration provided by the connective can first of all be viewed from an account in which readers are assumed to continuously predict discourse relations (Kehler, Kertz, Rohde, & Elman, 2008) (henceforth: the *relation prediction account*). Rohde and Horton (2014) show that readers can anticipate upcoming discourse relations, as revealed by anticipatory eye movements in a relation~location mapping task. In addition, readers’ expectations about upcoming discourse relations can also explain their interpretation of pronouns, even before the discourse relation becomes apparent (Kehler et al., 2008), suggesting that they indeed engage in predictive processing at the discourse level.

Such an account assumes that readers aim to establish a discourse-structural representation of the text as early as possible (in line with Cozijn et al., 2011). The question, however, is whether such an assumption about establishing a *relation* is necessary. Possibly, readers simply use the connective to make more specific predictions of the upcoming content (henceforth: *content prediction account*). In this scenario, the benefit provided by the connective is due to the connective enabling readers to more accurately predict the semantic content of the following clause. Consider (1): a reader is more likely to predict that Meghan is going to a party after encountering the connective *even so* than when the connective is not present. Evidence that readers adjust their expectations about the upcoming event based on connectives comes from various EEG studies (Xiang & Kuperberg, 2015; Köhne-Fuetterer, Drenhaus, Delogu, & Demberg, 2021). Thus, the *content prediction account* suggests that there is enhanced *semantic* predictability in the presence of a connective, facilitating processing.

Both the *relation prediction account* and the *content prediction account* can explain the facilitative effect of the connective, although they differ in the pro-

cessing levels at which they hypothesize readers to make predictions. A relation prediction account suggests that predictions are made at a discourse structural level, while a content prediction account suggests that predictions are made at a semantic level. Note that these two accounts are not mutually exclusive: It is well possible that the facilitative effect of the connective is due to both enhanced *relation* and *semantic* prediction. Nevertheless, disentangling these two accounts allows us to gain more insight into the processing instructions that connectives provide. We propose a modelling approach that allows us to make specific predictions on the role of the connective within a given context, formalizing the relation between expectations and processing difficulty using information theory. Surprisal, which denotes the expectancy of a linguistic signal given the context, is assumed to be proportional to processing difficulty (Levy, 2008). For example, higher surprisal has been shown to lead to longer reading times (Demberg & Keller, 2008; Wilcox, Pimentel, Meister, Cotterell, & Levy, 2023). In terms of the two accounts outlined above, the connective is assumed to reduce the uncertainty (i.e. surprisal) about the upcoming material¹ and thus reduces processing effort.

The relation prediction account hypothesizes that the benefit of a connective is proportional to how (un)expected the relation is without the connective, with a larger connective benefit when the relation is unexpected. The content prediction account suggests that the benefit of the connective crucially depends on whether the same content is more predictable with than without the connective. Although these types of predictability are correlated, there are contexts in which they diverge, as outlined in the next section. In the next section, we will specify how the two accounts can be modeled, followed by a description of the pretest from which the model predictions were derived. We then test the predictions of the model using a self-paced reading study. Our research questions are:

- Does semantic predictability explain the processing benefit of the connective, as posited by a content prediction account?
- Does relational predictability explain (additional) processing benefit, as posited by a relation prediction account?

The contribution of this paper is three-fold. First, we propose a cognitive model of the connective’s facilitative effect that allows us to make item-level predictions. Second, we test the predictions of two processing accounts, providing insight into the level of predictions that readers make. Finally, we show that the facilitative effect of the connective goes beyond providing more accurate relation and content prediction.

¹The two accounts differ in what this material consists of: the discourse relation as in relation prediction account or the content of the segment following the connective in content prediction account. This will be formalized in the next section.

Modeling the two accounts

Modeling the relation prediction account and the content prediction account allows us to make fine-grained predictions of the facilitative effect of the connective in different contexts, which enables us to test the validity of the two accounts. The relation prediction account assumes that the facilitative effect of the connective depends on the predictability of the relation. This predictability will be modeled as the surprisal of the relation and referred to as RP:

$$RP = \text{Surp}(rel) = -\log_2 p(rel|context) \quad (1)$$

The processing gain provided by a connective is then defined as the difference in this surprisal in the presence and absence of a connective. If a connective is not ambiguous (i.e. there is only a single relation that can plausibly follow the connective), surprisal will be 0. Gain is thus generally positive.

Obtaining a semantic-level measure of predictability is less straightforward. Surprisal estimates, as obtained from large language models (LLMs), conflate various aspects of predictability and therefore do not allow us to distinguish whether the target is predictable at the lexical, syntactic or semantic level. To isolate semantic predictability, we consider semantic information value (Giulianelli, Wallbridge, & Fernández, 2023) as our measure of content predictability (henceforth: CP). Information value is a measure that quantifies predictability of an utterance y given a context x as its distance d to plausible alternatives ($A_{context}$):

$$CP = IV(y|context) = d(y, A_{context}) \quad (2)$$

By considering how close the meaning of an utterance is from that of plausible alternatives, this measure captures that the content is semantically predictable if other likely continuations carry a similar meaning. Information value has been shown to predict processing difficulty similarly as and complementary to surprisal (Giulianelli et al., 2023). Again, the facilitative effect of the connective can be modeled by subtracting the information value in the explicit condition from that in the implicit condition.

Model predictions

Since a high content predictability entails a high relation predictability, the two accounts often make similar predictions about the extent to which a connective should facilitate reading. However, the content prediction account and relation prediction account yield different predictions with respect to how the size of the facilitative effect of the connective depends on the predictability of the *content*. Consider the two contexts in 2. The content of the clause following the material (in bold) is more predictable after reading 2a, than after 2b.

- (2) Angela used to live in a small flat in Atlanta.
 - a. She didn’t pay rent for months. high
 - b. She had over fifteen cats. low




Figure 1: Model predictions across semantic predictability and relation marking for the different accounts. The model predictions are hypothesized to be proportional to processing difficulty (i.e. reading times).

(Therefore,) she **was evicted** by her landlord.

A content prediction account hypothesizes that the effect of the connective should be smaller when the specific content is *not* predictable (i.e. in 2b), since in these cases the connective provides little information about the upcoming content. The content will remain (relatively) unpredictable, regardless of the presence of the connective. To illustrate the predictions of this model, we calculate the various measures of predictability for our experimental items, based on a pretest (see the next section). They are split by high or low semantic predictability and follow contexts with and without a connective.² These predictions are hypothesized to be proportional to processing difficulty, as measured by reading time. As can be seen in Figure 1a, the content prediction account predicts that, solely based on semantic predictability, utterances should be read faster in the presence than in the absence of a connective, which is a significant difference in the items in the present study ($t(47) = -4.26, p < .001$). In addition, the content prediction account assumes a main effect of semantic predictability, with longer reading times when CP is higher. As outlined above, the model crucially predicts that the facilitative effect of the connective should be smaller in contexts in which the content is less predictable.

A relation prediction account, on the other hand, leads to different hypotheses. Similar to the content prediction account, the model significantly predicts higher processing difficulty for implicit relations ($t(47) = -5.52, p < .001$). However, it assumes that the connective facilitates processing to the extent that it provides information about the upcoming *relation* rather than content. In other words, when the relation is easy to predict in the absence of a connective, there is little room for facilitation. In contexts where the relation is difficult to predict without a connective, the connective reduces the surprisal of the relation a lot and a larger facilitative effect of the connective is predicted. Although relation surprisal and semantic information

value are correlated in our pretest ($r = .31, CI = [.11, .48], p < .01$),³ the two accounts make opposite predictions about how this interacts with the facilitative effect of the connective: a relation prediction account hypothesizes an interaction with a larger facilitative effect of the connective in contexts for which the content is *unpredictable*, such as 2b, whereas a content prediction account predicts that this effect is larger for *predictable* contents.

Note that this version of the relation prediction account does not take into account that more surprising events are more difficult to process in itself, as illustrated by the lack of a difference between high and low semantic predictability in the explicit condition. It is unlikely, however, that semantic predictability does not contribute to processing difficulty at all: in addition to processing a relation, the reader also needs to process content. We will therefore also consider a combined account, which hypothesizes that both relational and semantic predictability influence processing, illustrated in Figure 1c for a scenario in which both types of predictability equally contribute to processing difficulty. Such a model does take into account that semantic predictability affects reading times, but also assumes that this is influenced by relational predictability. We remain agnostic with respect to the extent to which each type of predictability contributes to processing, but consider a significant effect of relational predictability on top of semantic predictability as evidence for a weak version of the relation prediction account. Evidence for the content prediction account is found if semantic predictability, as estimated with semantic information value, can explain the processing benefit of the connective alone. Evidence for the strong version of the relation prediction account is found if relational predictability, as estimated by relation surprisal, can explain the processing benefit of the connective alone.

Obtaining model estimates

We conducted a story continuation pretest to calculate the surprisal of relation (RP) and semantic information value (CP).

³This does not affect the VIF in our regression models later.

Participants 160 native English speakers (mean age: 39 years, 85 female), recruited via Prolific, participated.

Materials The materials consisted of 24 items with the following structure: (i) an introductory sentence, (ii) a sentence that was the first argument of the relation, (iii) a pronoun and (iv) a target region that was only present in the reading study, not this pretest. Each item consisted of four different versions, following a 2x2 design. All versions within an item ended with the same target region, but differed in (a) whether the pronoun was preceded by a connective (i.e. explicit) or not (i.e. implicit), and (b) whether the target region was either highly predictable or less predictable depending on the context. The latter distinction was based on an earlier pretest and was implemented to ensure variation in the semantic predictability of our items. Note that the manipulation of the predictability of the content also ensured variability in the predictability of the discourse relation, since these measures of predictability are (weakly) correlated, as discussed above. Also note that the goal of the current pretest is to obtain a continuous (rather than binary) measure of *semantic* and *relation* predictability that is estimated for each condition and context separately. To illustrate, an example of the different contexts for an item, along with the corresponding predictability values (CP | RP), is shown in 3 and 4 below.

- (3) Angela used to live in a small flat in Atlanta. She didn't pay rent for months.
 - a. Therefore, she ... (0.57 | 0.00)
 - b. She ... (0.72 | 0.23)
- (4) Angela used to live in a small flat in Atlanta. She had over fifteen cats.
 - a. Therefore, she ... (1.21 | 0.00)
 - b. She ... (1.19 | 1.15)

Procedure The items were distributed across 8 lists, each containing 12 experimental items (3 in each condition), as well as 18 fillers (6 implicit relations, 6 relations with “then” and 6 with “because”). Every list was completed by 20 participants, who provided a logical continuation to the prompt.

Analysis The continuations were analyzed for their relation and the event. Relation annotation was done by the first author, according to the PDTB3 guidelines.⁴ All items in the explicit condition were coded as ‘result’ relations. Relation surprisal was obtained by taking the log probability of the target relation per item per condition.

To estimate semantic predictability, we followed Giulianelli et al. (2023)’s approach. Continuations were first cleaned by removing typo’s and repeated pronouns. Subsequently, sentence embeddings were obtained for each continuation, as well as the target event, using Sentence-BERT (Reimers & Gurevych, 2019). For every item, we took

⁴The features +BELIEF and +SPEECHACT were excluded.




Figure 2: By-item gain predictions according to the relation prediction account and the content prediction account.

the mean Euclidian distance between the continuations and the target region per condition.

For each of the two accounts, the pretest yields four predicted values per item: one for each of the 2x2 conditions. For the content prediction account, these values consist of the different semantic information values (CP) for every of the four contexts. Likewise, for the strong version of the relation prediction account, this consists of the different relation surprisal values (RP). For the combined account, both the SP and the RP are taken. Using these values, we can also predict how much a connective facilitates reading given a certain context. To obtain the predicted *gain* of the connective for each item and context, we subtract the predicted value for the explicit condition for that context from that of the predicted value in the implicit condition for the same context.

Descriptives As can be seen in Figure 2, the two accounts predict the size of the facilitative effect of connective to be different in the same context. For example, for the item in the top left corner, the content prediction account predicts one of the largest effects of the connective ($CP_{gain} = 0.27$), whereas the strong version of the relation prediction account predicts a small effect ($RP_{gain} = 0.32$), compared to other items. Compare this with the item on the bottom right for which the relation prediction account predicts the largest benefit of the connective ($RP_{gain} = 4.32$), whereas the content prediction account even predicts a negative effect ($CP_{gain} = -.07$). We assume that this predicted negative effect is due to a lack of facilitation since there is higher predictability of the content without the connective, but theoretically it could also reflect a slow-down due to disconfirmed prediction.

Reading study

To examine which of the models best predicts the processing benefit provided by the connective, we conducted a self-paced reading study, examining the extent to which the connective facilitates reading the region directly following it.

Methodology

Participants 121 adult native speakers of English with no known reading disorder (mean age: 37 years; 54 female) participated in the experiment.

Materials and design Experimental materials consisted of the 24 items that were included in the model prediction estimate study. A spillover region as well as a spillover sentence were added to prevent any sentence or story wrap-up effects on the reading times of the area of interest, as in (5).

- (5) Angela used to live / in a small flat / in Atlanta. / She didn't pay rent / for months. / She / **was evicted** / by her landlord. / Angela decided / to move to a rural area.

Reading times were measured on the target region (in bold). The items were divided over 4 lists, with each list containing 6 items in every condition, and interspersed with 28 fillers.

Procedure Participants read the items in a non-cumulative moving window self-paced paradigm, implemented in PCIBex (Schwarz & Zehr, 2021). The order of the items was pseudo-randomized for every participant. Half of the items were followed by verification statements.

Analysis Data from participants who answered less than 70% of the verification statements correctly were removed from analysis ($n=8$). In addition, we removed data from trials in which participants spent more than a minute ($n=6$), as this indicates that they took a break from the experiment. Furthermore, we removed reading times on the target region above 2000 ms ($n=40$) or below 100 ms ($n=2$), as well as reading times that were more than 2.5 SD away from the participant's mean ($n=78$), removing 4.1% of the data points for the target region. 2782 data points were left for analysis.

Reading times of the target region were analyzed using mixed-effects regression models with a gamma identity link, using the `lme4` and `lmerTest` package (Bates, Mächler, Bolker, & Walker, 2015; Kuznetsova, Brockhoff, & Christensen, 2017) in R (R Development Core Team, 2008). In all models reported below, context, trial number, region position, length in characters and summed log word frequency were included as covariates. We aimed for a maximal random effect structure (Barr, Levy, Scheepers, & Tily, 2013).

Results

As can be seen from Figure 3, reading times of the target region are lower in the explicit condition compared to the implicit condition. Relation marking is indeed a significant predictor in the model (see Table 1). Figure 3 also suggests that the effect of relation marking is marginally smaller in contexts that yield low semantic predictability than in the highly predictable condition. This is contrary to the predictions of the strong version of the relation prediction account, illustrated in Figure 1b. To formally test the predictions of the different accounts, we examine whether relation marking still significantly predicts reading times when




Figure 3: Log-transformed raw reading times (bars: standard error) across semantic predictability and relation marking.

the model predictions of the respective accounts are included (cf. Delogu, Crocker, & Drenhaus, 2017). To illustrate, the content prediction account predicts that the facilitative effect of the connective can fully be attributed to differences in content predictability (CP) in the presence or absence of a connective. If this is true, the effect of connective should disappear when CP is accounted for, as all its variance can also (and more accurately) be explained by CP. In other words, relation marking should not significantly predict reading times when CP is also included as a predictor in the model.

As shown in Table 1, CP strongly and significantly predicts reading times overall, both when including this as the only measure of predictability (content), as well as when also taking into account RP (combined). The target region is read slower when its content is less predictable (i.e. the semantic information value is higher). Unlike CP, RP does not significantly predict reading times, suggesting that relation surprisal does not affect processing of the target region.

Crucially, however, there is still a significant facilitative effect of the connective when also including CP. Contrary to the predictions of the content prediction account, this indicates that CP does not account fully for the differences in reading times associated with the connective. The same holds when RP is added to the model, suggesting that relation surprisal alone can also not account for the facilitative effect of the connective. Even when adding both CP and RP to the model, as suggested by a combined account, there is still additional variance associated with relation marking left.

Although the approach above tests the strong version of the various accounts, namely whether semantic and/or relation predictability *fully* explains the facilitative effect of the connective, it does not show whether the model predictions of the various accounts *partly* explain the difference in reading times between explicit and implicit items. We therefore also examine whether these two accounts can predict the processing gain provided by the connective (i.e. the difference in reading times between the explicit and implicit condition), by first residualizing the reading times⁵ and then subtracting the

⁵base model: $rt \sim \text{context} + \text{trial} + \text{length} + \text{position} + \text{frequency} + (1 + \text{context} || \text{item}) + (1 + \text{context} | \text{ptcp})$

Table 1: Estimates for the predictors of interest per model.
***: $p < .001$, **: $p < .01$, *: $p < .05$

Model	Predictor	β	SE	t	p
base	connective	12.60	4.49	2.81	**
content	connective	11.97	4.30	2.78	**
	CP	32.35	8.73	3.71	***
relation	connective	10.15	4.25	2.39	*
	RP	5.75	4.53	1.27	.20
combined	connective	9.75	4.43	2.20	*
	CP	32.68	9.70	3.37	***
	RP	4.93	4.48	1.10	.27

Table 2: Model estimates of the predictors of interest for the residualized difference in reading times.

Model	Predictor	β	SE	t	p
content	CP_{gain}	-0.07	0.11	-0.63	.53
relation	RP_{gain}	0.01	0.01	1.07	.29
combined	CP_{gain}	-0.04	0.11	-0.36	.72
	RP_{gain}	0.01	0.01	0.92	.36

mean residualized reading time in the explicit condition from that in the implicit condition per item and item context. This difference in reading time is then regressed on the model gain predictions of the two accounts shown in Figure 2. As can be seen in Table 2 neither of the gain predictors was found to have a significant effect on the between-condition reading times, neither by itself nor combined. This suggests that neither semantic nor relation predictability predicts the facilitative effect of the connective.

Discussion

In the present study, we operationalized two types of predictability that have been hypothesized in the literature to influence processing: relation predictability (Kehler et al., 2008; Rohde & Horton, 2014) as the surprisal of the discourse relation and content predictability as semantic information value (Xiang & Kuperberg, 2015; Köhne-Fuetterer et al., 2021). In particular, we examined whether enhanced predictability in the presence of a connective can explain the processing benefit of the connective. Our results confirm previous findings that connectives facilitate processing of the material directly following it. We also show that semantic information value affects processing, in line with findings regarding predictability effects. Relation predictability did not significantly affect reading times of the target region. We thus do not find any evidence that readers make discourse-structural level predictions during normal reading. Crucially, neither semantic nor discourse relation predictability accounted for the facilitative effect of the connective. Thus, we do not find any evidence in favor of the content prediction account or the relation prediction account. The function of the connective seems to go above and beyond signalling to the reader what content, or what specific relation will follow.

Content predictability, which was operationalized as semantic information value, was found to strongly predict reading times. This is in line with previous research that predictable material, as indexed by semantic information value (Giulianelli et al., 2023) or surprisal estimates from LLMs (Wilcox et al., 2023), is easier to process than material that is less predictable and shows that readers make semantic-level predictions. Nevertheless, content predictability did not explain the difference in reading times caused by the presence of the connective.

It remains an open question what then constitutes the facilitative effect of a connective on processing. Our findings suggest that connectives provide more processing instructions than only informing the reader how the upcoming clause is related to the previous one (cf. Cozijn et al., 2011) or updating predictions on upcoming content (cf. Xiang & Kuperberg, 2015). Possibly, upon seeing a connective, readers adapt their processing strategy such that they *shallowly* process the material following the connective, only to process the relation more *deeply* at sentence wrap-up. This could explain why the effect of the connective seems to be unrelated to improved semantic or relation predictability. Previous studies have indeed shown that readers speed up immediately after encountering a connective, but that reading is often slowed down sentence-finally in explicit relations (Cozijn et al., 2011; van Silfhout et al., 2015). Nevertheless, this does not suggest that no semantic processing occurs clause-initially at all: implausible connectives have been shown to immediately disrupt processing (Canestrelli, Mak, & Sanders, 2013) and our own findings show that semantic predictability influences sentence-initial reading regardless of the presence of a connective.

Self-paced reading measures conflate various processes, in which connectives and surprisal might have differential effects. For example, connectives are known to trigger more, but shorter, regressions (van Silfhout et al., 2015), and surprisal has been shown to affect mostly early measures of processing (de Varda, Marelli, & Amenta, 2023). In future work, we will therefore use eye-tracking-while-reading to further explore in what way connectives provide processing instructions. In addition, we will examine whether other levels of predictability influence reading, and whether they interacts with the presence of a connective.

Conclusion

In line with previous research, the present study provides evidence that discourse processing is guided by semantic-level predictions and that discourse connectives facilitate sentence-initial reading. However, this effect of the connective cannot be attributed to enhanced semantic or relation predictability. Our findings thus suggest that the connective facilitates processing above and beyond informing the reader about the upcoming content, or even relation, and possibly change readers' processing strategy. In doing so, this line of research sheds further light on the role (and limits) of prediction in language processing.

Acknowledgments

This project is supported by the German Research Foundation (DFG) under Grant SFB 1102 (“Information Density and Linguistic Encoding”, Project-ID 232722074).

References

- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. doi: 10.1016/j.jml.2012.11.001
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. doi: 10.18637/jss.v067.i01
- Britton, B. K. (1994). *Understanding expository text: Building mental structures to induce insights*. Academic Press.
- Canestrelli, A. R., Mak, W. M., & Sanders, T. J. M. (2013). Causal connectives in discourse processing: How differences in subjectivity are reflected in eye movements. *Language and Cognitive Processes*, 28(9), 1394–1413.
- Cozijn, R., Noordman, L. G., & Vonk, W. (2011). Propositional integration and world-knowledge inference: Processes in understanding ‘because’ sentences. *Discourse Processes*, 48(7), 475–500.
- Delogu, F., Crocker, M. W., & Drenhaus, H. (2017). Teasing apart coercion and surprisal: Evidence from eye-movements and erps. *Cognition*, 161, 46–59.
- Demberg, V., & Keller, F. (2008). Data from eye-tracking corpora as evidence for theories of syntactic processing complexity. *Cognition*, 109(2), 193–210.
- de Varda, A. G., Marelli, M., & Amenta, S. (2023). Cloze probability, predictability ratings, and computational estimates for 205 english sentences, aligned with existing eeg and reading time data. *Behavior Research Methods*, 1–24.
- Gernsbacher, M. A. (1997). Coherence cues mapping during comprehension. *Processing interclausal relationships. Studies in the production and comprehension of text*, 3–22.
- Giulianelli, M., Wallbridge, S., & Fernández, R. (2023, December). Information value: Measuring utterance predictability as distance from plausible alternatives. In H. Bouamor, J. Pino, & K. Bali (Eds.), *Proceedings of the 2023 conference on empirical methods in natural language processing* (pp. 5633–5653). Singapore: Association for Computational Linguistics. Retrieved from <https://aclanthology.org/2023.emnlp-main.343> doi: 10.18653/v1/2023.emnlp-main.343
- Kehler, A., Kertz, L., Rohde, H., & Elman, J. L. (2008). Coherence and coreference revisited. *Journal of Semantics*, 25(1), 1–44.
- Köhne-Fuetterer, J., Drenhaus, H., Delogu, F., & Demberg, V. (2021). The online processing of causal and concessive discourse connectives. *Linguistics*, 59(2), 417–448. doi: 10.1515/ling-2021-0011
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13), 1–26. doi: 10.18637/jss.v082.i13
- Levy, R. P. (2008). Expectation-based syntactic comprehension. *Cognition*, 106(3), 1126–1177.
- R Development Core Team. (2008). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <http://www.R-project.org> (ISBN 3-900051-07-0)
- Reimers, N., & Gurevych, I. (2019, 11). Sentencebert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 conference on empirical methods in natural language processing*. Association for Computational Linguistics. Retrieved from <https://arxiv.org/abs/1908.10084>
- Rohde, H., & Horton, W. S. (2014). Anticipatory looks reveal expectations about discourse relations. *Cognition*, 133(3), 667–691. doi: 10.1016/j.cognition.2014.08.012
- Sanders, T. J. M., & Noordman, L. G. M. (2000). The role of coherence relations and their linguistic markers in text processing. *Discourse Processes*, 29(1), 37–60.
- Schwarz, F., & Zehr, J. (2021). Tutorial: Introduction to pcibex—an open-science platform for online experiments: Design, data-collection and code-sharing. In *Proceedings of the annual meeting of the cognitive science society* (p. 15-16). virtual.
- van Silfhout, G., Evers-Vermeul, J., & Sanders, T. J. M. (2015). Connectives as processing signals: How students benefit in processing narrative and expository texts. *Discourse Processes*, 52(1), 47–76.
- Wilcox, E. G., Pimentel, T., Meister, C., Cotterell, R., & Levy, R. P. (2023). Testing the predictions of surprisal theory in 11 languages. *Transactions of the Association for Computational Linguistics*, 11, 1451–1470.
- Xiang, M., & Kuperberg, G. (2015). Reversing expectations during discourse comprehension. *Language, Cognition and Neuroscience*, 30(6), 648–672. doi: 10.1080/23273798.2014.995679