Syntactic adaptation to short-term cue-based distributional regularities

Anonymous CogSci submission

Abstract

Syntactic adaptation to short-term exposure has been documented with both single-trial priming and cumulative priming paradigms. These studies usually involve repeated exposure to the same structure (e.g. reduced relative clauses), and therefore it remains open whether people can track *context-dependent* regularities through short-term exposure. In the current study, we present a self-paced-reading experiment that investigates context-dependent syntactic adaptation by manipulating the relationship between the animacy feature of the subject NP (animate vs. inanimate) and the corresponding parse of a verb following a subject NP. We analyze the results in terms of a log-linear model for context-dependent syntactic adaptation. The results suggest that comprehenders can track and adapt to cue-based distributional regularities, but only when the short-term regularities are consistent with the long-term ones existent in their native language.

Keywords: Syntactic adaptation; syntactic priming; context-dependent adaptation; cue-based regularities; statistical learning; log-linear models; GPT-3 language model

Introduction

Since the seminal work of Bock (1986), it is widely observed that speakers tend to reuse the same structure they just experienced, a phenomenon known as **syntactic priming**. The basic finding has been consistently replicated in other studies using different syntactic constructions (Bock, 1986; Scheepers, 2003; Mahowald, James, Futrell, & Gibson, 2016), across languages (Hartsuiker & Kolk, 1998; Shin & Christianson, 2009), and in bilingual contexts (Hartsuiker, Pickering, & Veltkamp, 2004). In comprehension, it is also found that comprehenders adjust their expectations on the possible parses of a string depending on what structure they have parsed earlier (Pickering & Ferreira, 2008; Tooley & Traxler, 2010). Syntactic priming is a form of linguistic adaptive behavior to short-term exposure. It is of interest to psycholinguists for a number of reasons, one of which is that understanding syntactic priming provides insights into the mechanism of implicit learning.

Using a cumulative priming paradigm, Fine, Jaeger, Farmer, and Qian (2013) aim to connect syntactic priming to statistical learning theory (Saffran, Aslin, & Newport, 1996). The idea is that syntactic priming results from a short-term shift in the comprehender's model of the statistical regularities in the linguistic input. Because the process of online language comprehension is related to the incremental predictability of linguistic objects (Garnsey, Pearlmutter, Myers, & Lotocky, 1997; MacDonald, Pearlmutter, & Seidenberg, 1994; Trueswell, 1996; Reali & Christiansen, 2007; Hale, 2001; Levy, 2008), a shift in the comprehender's probability model will result in characteristic changes in language processing as reflected in dependent measures such as reading time.

However, previous studies of syntactic adaptation, either based on the trial-to-trial or the cumulative priming paradigm, have been limited in that they usually involve only repeated exposure to the same structure (for example, repeated exposure to reduced-relative clauses), regardless of any other sources of linguistic information. The exposure usually modulates the frequency of a single structure on its own, and thus effectively raises or lowers the probability of that structure in the comprehender's expectations independent of context. However, language processing is known to be highly context-dependent: the processing of a specific structure depends on more than just its frequency out of context, but is also conditioned on fine-grained contextual cues (MacDonald et al., 1994). For example, expectations for a reduced relative clause versus a main verb are modulated by the animacy of the preceding noun (Trueswell, Tanenhaus, & Garnsey, 1994).

In this work, through behavioral experiments and computational modeling, we investigate how short-term adaptation can modulate the comprehender's expectations about the statistical *relationship* between contextual cues and syntactic structures. The specific syntactic structure is the reduced-relative clause garden-path sentences. As shown in (1) below (Fine et al., 2013), the verb "warned" is temporally ambiguous between a reduced relative (RR) interpretation (1a) and a main verb (MV) interpretation (1b). Due to the high frequency of the MV structure in English, upon encountering this ambiguous verb, the parser would favor the MV interpretation. If the sentence turns out to be an RR structure as in (1a), there will be increased reading time at the disambiguating main verb "conducted" in (1a). This is known as the garden-path ambiguity effect (Frazier & Fodor, 1978).

(1) The experienced soldiers...

  a ...warned about the dangers conducted the midnight raid.

  b ...warned about the dangers before the midnight raid.

Through the lens of the MV/RR ambiguity, we study the relationship between the animacy feature of a subject NP (Animate vs. Inanimate) and the parse of the locally ambiguous subject-verb combination – whether it is interpreted as a MV or RR structure before disambiguation.

In the remainder of this paper, we first present a self-paced reading (SPR) experiment whose results suggest that comprehenders indeed track and adapt to cue-based context-dependent short-term regularities. But such adaptations only take place when the short-term regularities are consistent with people's long-term knowledge. In the second part of the paper, we will present a log-linear model for the context-dependent adaptation effect we observed.

Experiment

In an SPR experiment, we examined whether and how the garden-path ambiguity of a reduced-relative vs. a main-verb parse changes after short-term exposure to linguistic materials in which the relationship between the parse of a verb and the animacy of the dependent noun was manipulated. The experiment consists of two blocks: an exposure block and a testing block. Block 1 (the exposure block) manipulated the co-occurrence statistics between the animacy feature on the subject-NP and the MV/RR parse on the verb following the subject-NP. The experiment was a between-subject design with three treatment groups. The differences between these groups lie in the exposure block. As summarized in Table 1, in the exposure block, **Group A** read 20 RR sentences with animate subjects and 20 MVs with inanimate subjects—this mapping is inconsistent with the long-term regularities present in English, where RR sentences typically come with an inanimate subject. In contrast, **Group B** read 20 RRs with inanimate subjects and 20 MVs with animate subjects. Group B's exposure is consistent with their general long-term experience. **Group C** is the control group: participants read 40 filler sentences in the exposure. Sample stimuli for Block 1 are presented in (2).

Block 2 (the testing block) is identical across all three groups, where all the participants read 8 ambiguous RRs and 8 unambiguous RCs (4 animate and 4 inanimate for both sentence types). The sample stimuli of Block 2 are presented in (3), where the critical disambiguating region and the spill-over region are underlined. The verbs used in Block 2 also appeared in Block 1.

(2) Sample stimuli in Block 1 (the exposure block)

- Group A

  The defendant/ examined/ by the lawyer/ turned out/ to be/ unreliable. [animate, RR]

  The hypothesis/ examined/ the factors/ that/ affected/ hearing. [inanimate, MV]

- Group B

  The hypothesis/ examined/ by the scientist/ was not/ widely known. [inanimate, RR]

  The defendant/ examined/ the testimony/ carefully/ yesterday. [animate, MV]

(3) Sample stimuli in block 2 (the testing block)

- The patient/ examined/ by the doctor/ was diagnosed/ with diabetes. [animate, ambiguous]
- The patient/ that/ was/ examined/ by the doctor/ was diagnosed/ with diabetes. [animate, unambiguous]
- The document/ examined/ by the lawyer/ turned out/ to be/ unreliable. [inanimate, ambiguous]
- The document/ that/ was/ examined/ by the lawyer/ turned out/ to be/ unreliable. [inanimate, unambiguous]

Predictions If comprehenders can adapt syntactic expectations towards short-term cue-based distributional regularities, we expect the garden-path ambiguity effect of Group A to be weakened for animate subjects and be enhanced for inanimate subjects. In contrast, the garden-path effect of Group B is expected to be enhanced for animate subjects and be weakened for inanimate subjects. Moreover, according to the **inverse frequency effect** observed in previous studies (Kaschak, Kutta, & Jones, 2011; Reitter, Keller, & Moore, 2011), less frequent structure should induce larger learning error, resulting in stronger adaptation effect. In our design, the animacy–parse relationship in Group A is in an opposite direction to the long-term knowledge of English speakers. This is expected to be associated with larger learning error compared to Group B, whose animacy–parse relationship is consistent with the long-term pattern in English. Therefore, we also expect the magnitude of the adaptation effect in Group A to be larger than Group B.

Method

Participants and Procedure 400 native speakers of English living in the U.S. were recruited via Prolific (University of xxxx IRB18-0381). They were directed to PCIbex to take the experiment. Participants were randomly assigned to one of the three groups. For both blocks, in each trial, the participants read a sentence in the moving-window SPR paradigm: the sentence was first presented as a series of obscured word/phrase chunks and the participants pressed the space bar to reveal one word/phrase at a time. For some trials, the participants were asked to answer a comprehension question regarding the sentence they read. The comprehension questions were aimed at encouraging the participants to stay focused. There were eight practice trials before the main experiment session to familiarize participants with the SPR paradigm.

Results and discussion

Participants whose accuracy on comprehension questions was less than 80% were excluded, yielding a final data set with 122 participants in Group A, 126 in Group B, and 125 in Group C. Observations with a reading time lower than 100 ms or higher than 5000 ms were removed. We then further removed reading times that were more than three standard deviations away from the mean per region and per condition. The reading times of Block 2 (the testing block) by condition and by group are presented in Figure 1.

Figure 2 shows the average garden-path effect, defined as the reading time for the ambiguous condition minus the unambiguous condition across experimental conditions in the

Table 1: Experiment block design

Group	Block 1	Block 2
A	20 Ani-RRs, 20 Inani-MVs	4 Ani-RRs, 4 Inani-RRs, 4 Ani-RCs, 4 Inani-RCs, 16 fillers
B	20 Inani-RRs, 20 Ani-MVs	4 Ani-RRs, 4 Inani-RRs, 4 Ani-RCs, 4 Inani-RCs, 16 fillers
C	40 fillers	4 Ani-RRs, 4 Inani-RRs, 4 Ani-RCs, 4 Inani-RCs, 16 fillers

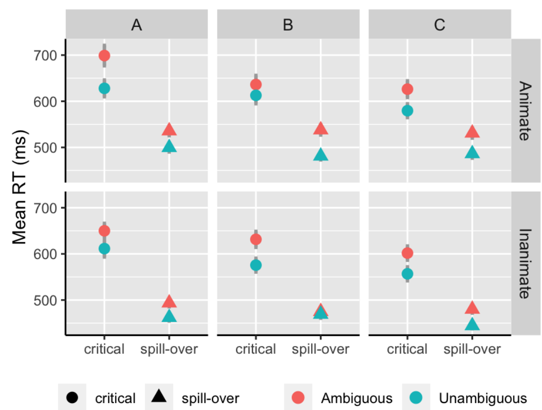

Figure 1: **Block 2** (testing block) reading times by group and by animacy. Effects appear in spill-over region of Group B.

spillover region. Visual inspection suggests that the garden-path effect for animate subjects, compared to the control Group C, is suppressed in Group A, whereas it is enhanced in Group B. For inanimate subjects, however, the effect mainly emerges in Group B with a substantially suppressed garden-path effect. As we will see, a statistical analysis partially confirms these observations.

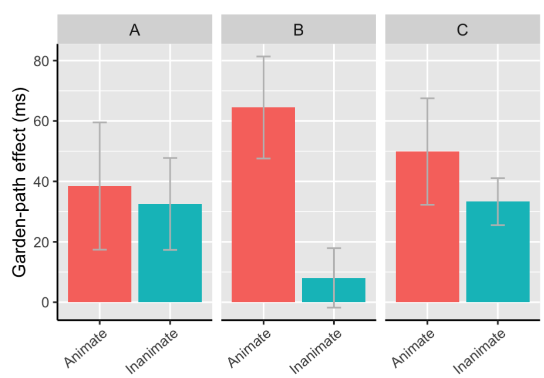

Figure 2: Human garden-path ambiguity effect (average RT for ambiguous trials minus unambiguous trials) in self-paced-reading experiment in the spill-over region by group.

We fit linear mixed-effect models to log-transformed RTs in the critical disambiguating and the spill-over regions of Block 2, with maximal random effects (Barr, Levy, Scheepers, & Tily, 2013). We focus on whether the garden-path ambiguity cued by subject animacy in Group A or B differs from the control Group C. For the B vs. C comparison, no effects were found on the disambiguating region. On the spill-over region, for Group B only, as in Model (4) below, there is an Animacy $\times$ Ambiguity interaction ($\hat{\beta} = 0.02$, $p < .01$): the ambiguity effect is significantly reduced for trials with inanimate but not animate subjects. No such interaction was found for Group C ($\hat{\beta} = 0.007$, $p = 0.33$). However, we did not find an Animacy $\times$ Ambiguity $\times$ Group interaction ($\hat{\beta} = 0.005$, $p = 0.24$) when the data from B and C are considered together as in Model (5). We speculate that the lack of three-way interaction could be due to the low statistical power. We conducted a post-hoc power analysis targeting the three-way interaction using SIMR package (Green & MacLeod, 2016) in R. The result shows that with a small effect size (approximately 20ms), 250 participants are needed for each group to achieve 80% power to detect the three-way interaction. For the A vs. C comparison, no relevant effect was detected on either the disambiguating or the spill-over region.

(4) Linear mixed-effect model for each individual group

logRT $\sim$ logRT.previous.region + Word.length + Ambiguity $*$ Animacy + (1 + Ambiguity $*$ Animacy | Subj) + (1 + Ambiguity | Item)

(5) Linear mixed-effect model with group contrasts for A vs. C and B vs. C

logRT $\sim$ logRT.previous.region + Word.length + Ambiguity $*$ Animacy $*$ Group + (1 + Ambiguity $*$ Animacy | Subj) + (1 + Ambiguity $*$ Group | Item)

The experimental result shows that participants can track and adapt to cue-based (animacy) context-dependent short-term regularities, but they only do so when short-term regularities in the exposure phase are consistent with their long-term knowledge (as shown in Group B). The fact that we did not observe a reliable adaptation effect in Group A, in which the regularities in the exposure phase are inconsistent with the long-term patterns in English, is unexpected from an error-driven learning mechanism that predicts inverse-frequency effect. Under this hypothesis, the more unexpected exposure should lead to larger learning effect, contrary to the findings from Group A.

Modeling

We present a log-linear model to capture the syntactic adaptation effect under the experimental setting introduced in the previous section. We generate the model-predicted surprisal of the RR parse, which we then compare with the garden-path ambiguity observed in our SPR experiment. The model allows us to quantitatively describe and estimate how short-term exposure in Groups A and B changes the statistical association strength between animacy and verb parse.

The model

Prior to training in the exposure phase, the probability of RR parse given that the subject NP is animate can be written in terms of (i) the overall bias toward the RR parse independent of context and (ii) the strength of association between animacy and RR parse, as in Eq. (1). Here, $w^{\mathrm{RR}}_{\mathrm{ani}}$ is the long-term strength of association between animate subject NP and RR parse and b^{RR} is the long-term bias towards RR parse.

$$
\begin{aligned}
p(\mathrm{RR} \mid c_{\mathrm{ani}}) &= \frac{e^{w^{\mathrm{RR}}_{\mathrm{ani}}+b^{\mathrm{RR}}}}{e^{w^{\mathrm{RR}}_{\mathrm{ani}}+b^{\mathrm{RR}}}+e^{w^{\mathrm{MV}}_{\mathrm{ani}}+b^{\mathrm{MV}}}} \\
&= \frac{1}{1+e^{(w^{\mathrm{MV}}_{\mathrm{ani}}-w^{\mathrm{RR}}_{\mathrm{ani}})+(b^{\mathrm{MV}}-b^{\mathrm{RR}})}}.
\end{aligned} \tag{1}
$$

We propose to represent the adaptation after the training by adding or subtracting an adaptation term k to the strength of the corresponding animacy-parse association. This allows us to study changes in association strength explicitly.[1] For group A, the animate subject is mapped to RR parse in the exposure phase, and the new weights w after adaptation is given by:

$$w^{\mathrm{RR}}_{\mathrm{ani}}(\mathrm{A}) = w^{\mathrm{RR}}_{\mathrm{ani}} + k_A \tag{2}$$

$$w^{\mathrm{RR}}_{\mathrm{inani}}(\mathrm{A}) = w^{\mathrm{RR}}_{\mathrm{inani}} - k_A \tag{3}$$

In contrast, for Group B, we increase the association of RR with animate subjects, and decrease the association with inanimate subjects:

$$w^{\mathrm{RR}}_{\mathrm{ani}}(\mathrm{B}) = w^{\mathrm{RR}}_{\mathrm{ani}} - k_B \tag{4}$$

$$w^{\mathrm{RR}}_{\mathrm{inani}}(\mathrm{B}) = w^{\mathrm{RR}}_{\mathrm{inani}} + k_B \tag{5}$$

Above, we described a log-linear model of the comprehender's expectations about the statistical regularities in language, and how these expectations change with exposure. In order to link this model to reading time data, we use **surprisal theory** (Levy, 2008; Hale, 2001), which holds that the reading time response at word w is proportional to the surprisal of word w given its preceding context c:

$$\mathrm{RT} \propto -\ln p(\mathrm{w} \mid c). \tag{6}$$

Under some assumptions about our MV/RR experimental materials, we can derive that the RT *difference* between the ambiguous and unambiguous conditions at the disambiguating word should be proportional to

$$-\ln p(\mathrm{RR} \mid c), \tag{7}$$

where c is the context for ambiguous condition.[2] In SPR data, this effect can be expected to appear in the spill-over region.

Using the log-linear model outlined above, the garden path effect (ambiguous minus unambiguous RT) before adaptation for an animate head noun should be given by:

$$
\begin{aligned}
\mathrm{RT\ effect} &\propto -\ln p(\mathrm{RR} \mid c_{\mathrm{ani}}) \\
&= \ln\left(1+e^{w'_{\mathrm{ani}}+b'}\right),
\end{aligned} \tag{8}
$$

where we have introduced the notation $w'_{\mathrm{ani}} \equiv w^{\mathrm{MV}}_{\mathrm{ani}} - w^{\mathrm{RR}}_{\mathrm{ani}}$ for the difference in the strength of association with animacy for the two parses, and analogously $b' \equiv b^{\mathrm{MV}} - b^{\mathrm{RR}}$ for the bias term. Eq. (8) holds similarly when the subject noun is inanimate, using w'_{inani}. After adaptation, the reading time response given animate subjects for Group A and B is given in Eq. (9) and (10) respectively;[3] the equations for inanimate subjects are given in (11) and (12). The higher the value of k, the larger the adaptation effect is. Since Group C received no training in the exposure phase, $k_C = 0$ and its reading time response should be the same as the long-term regularities represented in Eq. (8).

$$
\begin{aligned}
\mathrm{RT}(\mathrm{A} \mid c_{\mathrm{ani}})\ \mathrm{effect} &\propto -\ln p(\mathrm{RR} \mid c_{\mathrm{ani}}) \\
&= \ln(1+e^{((w'_{\mathrm{ani}}-k_A)+b'})
\end{aligned} \tag{9}
$$

$$
\begin{aligned}
\mathrm{RT}(\mathrm{B} \mid c_{\mathrm{ani}})\ \mathrm{effect} &\propto -\ln p(\mathrm{RR} \mid c_{\mathrm{ani}}) \\
&= \ln(1+e^{((w'_{\mathrm{ani}}+k_B)+b'})
\end{aligned} \tag{10}
$$

$$
\begin{aligned}
\mathrm{RT}(\mathrm{A} \mid c_{\mathrm{inani}})\ \mathrm{effect} &\propto -\ln p(\mathrm{RR} \mid c_{\mathrm{inani}}) \\
&= \ln(1+e^{((w'_{\mathrm{inani}}+k_A)+b'})
\end{aligned} \tag{11}
$$

$$
\begin{aligned}
\mathrm{RT}(\mathrm{B} \mid c_{\mathrm{inani}})\ \mathrm{effect} &\propto -\ln p(\mathrm{RR} \mid c_{\mathrm{inani}}) \\
&= \ln(1+e^{((w'_{\mathrm{inani}}-k_B)+b'})
\end{aligned} \tag{12}
$$

Simulation

Estimating parameters w' and b' We estimated the association strength w' and bias b' parameters for the log-linear model using a combination of corpus frequencies and the GPT-3 language model (Brown et al., 2020).

We first estimated the bias b' based on the structural probability of a reduced relative clause obtained from the Penn Treebank (Marcus, Marcinkiewicz, & Santorini, 1993), with $p(\mathrm{RR}) = 0.008$.[4] Based on this probability, the bias term was

[1] Our method differs from the syntactic adaptation model of van Schijndel and Linzen (2018), where models are fine-tuned to the exposure materials.

[2] The conclusion requires four assumptions: (1) there are at most two possible parses (MV and RR) for the preceding context c, (2) the probability of the disambiguating word w is zero under the MV parse, (3) the probability of the RR parse given the unambiguous context is 1, and (4) the probability of the disambiguating word given the RR parse is the same for both the ambiguous and unambiguous contexts.

[3] Although the proportion of RR parse in the short-term exposure phase of our experiment is much higher than the long-term linguistic environment of English, since the adaptation effect induced by the short-term regularities of a particular structure as a whole is beyond the scope of the current study, we did not apply an adaptation coefficient to the bias term b.

[4] The search query was `(NP-SBJ !<< @VP) $+ @VP` for MV construction (259298 occurrences), and `NP-SBJ < (NP $ @VP)` for RR parse (2101 occurrences).

Table 2: Estimates of w' and model-predicted surprisal of RR parse $h(\text{RR})$ by group

			A (k = -0.04)		B (k = 1.81)		C (k = 0)	
Verb	w'_{ani}	w'_{inani}	$h(\text{RR}\mid\text{ani})$	$h(\text{RR}\mid\text{inani})$	$h(\text{RR}\mid\text{ani})$	$h(\text{RR}\mid\text{inani})$	$h(\text{RR}\mid\text{ani})$	$h(\text{RR}\mid\text{inani})$
examine	-1.14	3.44	0.64	4.40	1.83	2.69	0.62	4.44
follow	1.92	1.24	2.99	2.29	4.72	0.92	2.95	2.33
find	–	0.57	–	1.71	–	0.57	0.90	1.74
recommend	1.36	–	2.47	–	4.17	–	2.44	-1.28
ask	1.62	-1.09	2.71	0.62	4.42	0.14	2.67	0.64
gather	–	-2.22	–	0.25	–	0.05	-1.21	0.26
report	4.72	0.60	5.74	1.73	7.51	0.59	5.70	1.77
analyze	0.86	0.55	2.03	1.70	3.68	0.57	2.00	1.73
accompany	–	0.51	–	1.67	–	0.55	-1.52	1.70
separate	1.74	3.27	2.82	4.23	4.54	2.53	2.78	4.27
fund	2.81	2.43	3.86	3.40	5.61	1.78	3.82	3.44
introduce	1.93	-1.02	3.00	0.66	4.73	0.15	2.96	0.67
provide	-0.81	-0.91	0.80	0.71	2.11	0.16	0.78	0.73
respect	7.38	4.13	8.41	5.08	10.18	3.34	8.73	5.12
support	1.44	-0.58	2.54	0.89	4.24	0.22	2.51	0.91
expect	0.74	2.31	1.93	3.29	3.57	1.69	1.90	3.33
Average	2.13	1.84	2.27	1.96	3.96	0.71	2.24	1.99

estimated by solving Eq. (13) for b'.

$$p(\text{RR}) = \frac{1}{1+e^{b'}}. \tag{13}$$

The animacy-parse association strength w' was estimated using the GPT-3 Language Model (LM). This value could not be estimated from corpus frequencies due to low frequency. We first generated a LM-predicted garden-path effect for each verb-animacy combination, and mapped it to the probability of RR structure given context $p(\text{RR} \mid c)$. To do this, we fed GPT-3 with sentences from the stimuli of the testing block in our experiment. For each sentence, we generated a suffix surprisal $h(\text{suffix})$: the preamble that goes before the critical *by-phrase* serves as the context c for GPT-3 model, and the suffix surprisal is the GPT-3 surprisal for the rest of the sentence given this context c, which was calculated by summing up the surprisal progressively generated at each word.[5] Sentences were paired in the sense that each pair corresponds to a verb-animacy combination and contains a locally ambiguous RR and its unambiguous RC counterpart, as in (3). For each pair, the LM-predicted garden-path effect was taken as the difference in suffix surprisal between ambiguous RR and unambiguous RC.[6] We then mapped this LM-predicted garden-path effect to the surprisal of RR structure given c, from which we obtained $p(\text{RR} \mid c)$:

$$\begin{aligned} -\ln p(\text{RR} \mid c) &= h(\text{RR} \mid c) \\ &= h(\text{suffix} \mid c_{\text{ambig}}) - h(\text{suffix} \mid c_{\text{unambig}}). \end{aligned} \tag{14}$$

Finally, we found the association strength values by solving for w'_{ani} (and analogously for w'_{inani}) given the corpus-estimated value of b' and the LM-estimated value of $p(\text{RR} \mid \text{animate})$ in the log-linear model:

$$p(\text{RR} \mid c_{\text{ani}}) \propto e^{w'_{\text{ani}}+b'}. \tag{15}$$

Adaptation coefficient k We estimated the adaptation coefficient k from the empirical data in our SPR experiment.[7] We first linked the LM-predicted garden-path effect to the empirical garden-path effect manifested on reading time through a conversion factor λ:

$$\text{RT effect} = -\lambda \ln p(\text{RR} \mid c). \tag{16}$$

The conversion factor λ was estimated from our empirical reading time data in the spill-over region of Group C using linear regression and the GPT-3 estimate for the probabilities $p(\text{RR} \mid c)$. The estimated λ is around 19.82. With this λ, we found the adaptation coefficients k for Group A and Group B separately by fitting Eqs. (9)–(12) to the average reading time

[5] We calculated the suffix surprisal over the entire rest of the sentence because the GPT-3 language model might not fully disambiguate the context at the by-phrase alone.

[6] This procedure gives negative LM-predicted garden-path effects for some verb–animacy combinations. Therefore, the average cue weights are obtained from the average LM-predicted garden-path effect, instead of being calculated by averaging cue weights estimated individually for each verb.

[7] Since the adaptation effect primarily showed up in the spill-over region, not the critical region, the estimation of k was only based on the empirical data in the spill-over region. For the same reason, later in the section of results and discussion, we only compared model-predicted garden-path effect with the empirical garden-path effect in the spill-over region.

data in the spill-over region using linear regression. We find $k_A = -0.04$ and $k_B = 1.81$, indicating a much larger change in association strength for Group B than Group A. It is worth noting that the adaptation for Group A learned by the model is not in the same direction as the empirical data: compared to Group C, the model predicts weakened cue strength for inanimate subjects, not for animate subjects. This is mainly because, before adaptation, the difference in LM-predicted garden-path effect between animate and inanimate subjects is smaller than the empirical pattern in Group C. In this sense, to fit the empirical garden-path effect in Group A, cue weights need to change in the opposite direction against what we expected. However, since the empirical data does not show significant interaction between animacy and garden-path effect either in Group A or in Group C, the influence of this unexpected but very small negative k on Group A should be limited.

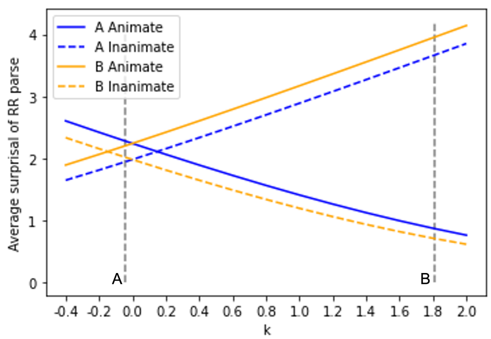

Figure 3: Average estimated surprisal with different k values. k values for Group A and B are marked by vertical dash lines.

We analyzed how the value of k influences the LM-predicted surprisal of RR. As seen in Figure 3, when k increases, for Group A, the surprisal of RR increases for inanimate subjects and decreases for animate subjects. That means, for Group A, animate subjects become more associated with RR parse, whereas inanimate subjects become more associated with MV parse, which is consistent with the short-term distributional regularities in the exposure phase of our experiment. Moreover, the association between animacy feature and MV/RR parses is reversed at around $k = 0.2$, beyond which an animate subject induces a smaller surprisal for the RR parse. For Group B, the trend is reversed, such that the surprisal of RR increases for animate subjects and decreases for inanimate subjects, indicative of a strengthened association between inanimate subjects and RR parse.

Results and discussion

For each verb, with the estimated w', b', and k presented in the previous section, we generated the model-predicted garden-path effect, which is the surprisal of RR parse, given animate and inanimate subjects respectively. The results are summarized in Table 2.[8] We also mapped the model-predicted garden-path effect to reading times based on the conversion factor λ estimated above, as visualized in Figure 4.

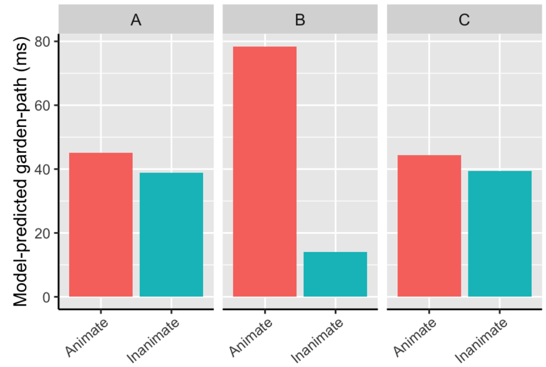

Figure 4: Model-predicted garden-path effect by group

The result shows that the pattern predicted by our log-linear model sufficiently fits the empirical garden-path effect shown in Figure 2. The simulated surprisal predicts a larger garden-path effect for animate condition than inanimate condition, showing that the model is able to capture the long-term distributional regularities in English, where the inanimate subject is more closely linked to RR interpretation, although such difference between animate and inanimate subjects is smaller than our SPR data as pointed out above. Similar to our empirical observation, the model is able to capture an adaptation effect. After the short-term exposure to sentences where inanimate subjects are exclusively linked to RR parse (Group B), the garden-path effect is reduced for inanimate subject and enhanced for animate subjects. The adaptation induced by the exposure to sentences where animate subjects are linked to RR parse (Group A) is limited, as shown in our SPR data. The absolute value of adaptation coefficient k for Group B ($k = 1.81$) is larger than for Group A ($k = -0.04$), further supporting the asymmetric adaptation effect magnitude between Group A and Group B.

Conclusions

To sum up, the current study shows that comprehenders can track and adapt to cue-based context-dependent short-term regularities, but possibly in a constrained fashion: we only found evidence for adaptation when the short-term regularities are consistent with participants' long-term knowledge. We also introduced a log-linear model for the context-dependent syntactic adaptation. The model quantitatively supports the asymmetric adaptation effect observable from the empirical data. Overall, this is indicative of a larger learning difficulty, at least at the cue-based level, for unexpected short-term regularities, and thus casts questions on proposals of error-driven learning based on previously observed inverse-frequency effect, which suggests that more unexpected exposure can lead to larger learning effect.

[8]The average model-predicted garden-path effect in Group A and B was obtained from the average LM-predicted garden-path effect before adaptation.

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.

Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive psychology*, *18*(3), 355–387.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... others (2020). Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.

Fine, A. B., Jaeger, T. F., Farmer, T. A., & Qian, T. (2013). Rapid expectation adaptation during syntactic comprehension. *PloS one*, *8*(10), e77661.

Frazier, L., & Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition*, *6*(4), 291–325.

Garnsey, S. M., Pearlmutter, N. J., Myers, E., & Lotocky, M. A. (1997). The contributions of verb bias and plausibility to the comprehension of temporarily ambiguous sentences. *Journal of memory and language*, *37*(1), 58–93.

Green, P., & MacLeod, C. J. (2016). Simr: an r package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution*, *7*(4), 493–498.

Hale, J. T. (2001). A probabilistic Earley parser as a psycholinguistic model. In *Proceedings of the second meeting of the north american chapter of the association for computational linguistics and language technologies* (pp. 1–8).

Hartsuiker, R. J., & Kolk, H. H. (1998). Syntactic persistence in dutch. *Language and Speech*, *41*(2), 143–184.

Hartsuiker, R. J., Pickering, M. J., & Veltkamp, E. (2004). Is syntax separate or shared between languages? cross-linguistic syntactic priming in spanish-english bilinguals. *Psychological science*, *15*(6), 409–414.

Kaschak, M. P., Kutta, T. J., & Jones, J. L. (2011). Structural priming as implicit learning: Cumulative priming effects and individual differences. *Psychonomic bulletin & review*, *18*(6), 1133–1139.

Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, *106*(3), 1126–1177.

MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review*, *101*(4), 676.

Mahowald, K., James, A., Futrell, R., & Gibson, E. (2016). A meta-analysis of syntactic priming in language production. *Journal of Memory and Language*, *91*, 5–27.

Marcus, M. P., Marcinkiewicz, M. A., & Santorini, B. (1993). Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, *19*(2), 313–330.

Pickering, M. J., & Ferreira, V. S. (2008). Structural priming: a critical review. *Psychological bulletin*, *134*(3), 427.

Reali, F., & Christiansen, M. H. (2007). Processing of relative clauses is made easier by frequency of occurrence. *Journal of Memory and Language*, *57*(1), 1–23.

Reitter, D., Keller, F., & Moore, J. D. (2011). A computational cognitive model of syntactic priming. *Cognitive science*, *35*(4), 587–637.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*(5294), 1926–1928.

Scheepers, C. (2003). Syntactic priming of relative clause attachments: Persistence of structural configuration in sentence production. *Cognition*, *89*(3), 179–205.

Shin, J.-A., & Christianson, K. (2009). Syntactic processing in korean–english bilingual production: Evidence from cross-linguistic structural priming. *Cognition*, *112*(1), 175–180.

Tooley, K. M., & Traxler, M. J. (2010). Syntactic priming effects in comprehension: A critical review. *Language and Linguistics Compass*, *4*(10), 925–937.

Trueswell, J. C. (1996). The role of lexical frequency in syntactic ambiguity resolution. *Journal of memory and language*, *35*(4), 566–585.

Trueswell, J. C., Tanenhaus, M. K., & Garnsey, S. (1994). Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language*, *33*, 285–318.

van Schijndel, M., & Linzen, T. (2018). A neural model of adaptation in reading. In *2018 conference on empirical methods in natural language processing (emnlp 2018).*