Syntactic Priming Depends on Procedural, Reward-Based Computations: Evidence from Experimental Data and a Computational Model

Yuxue Cher Yang (chery@uw.edu)
Department of Psychology, University of Washington
Campus Box 351525, Seattle, WA 98195 USA

Andrea Stocco (stocco@uw.edu)
Department of Psychology and Institute for Learning and Brain Sciences (I-LABS), University of Washington
Campus Box 351525, Seattle, WA 98195 USA

Abstract

Syntactic priming (SP) is the effect by which, in a dialogue, the current speaker tends to re-use the syntactic constructs of the previous speakers. SP has been used as window into the nature of syntactic representations within and across languages. Because of its importance, it is crucial to understand the mechanisms behind it. Currently, two competing theories exist. According to the surprisal theory, SP is driven by the mismatch with internal predictions and enhanced by factors that enhance surprise (i.e., use of low-frequency verbs). According to the declarative theory, SP is driven by the re-activation of declarative memory structures that encode template structures. Here, we propose a third and novel hypothesis, namely, that SP is driven by the successful application of procedural knowledge, in agreement with Ullman's model. This hypothesis makes the unique prediction that SP will be reversed when the prime sentence includes grammatical errors, but not semantic errors. The theory is supported by a computational model. An experiment confirmed the prediction of the theory.

Keywords: Syntactic Priming, Procedural Knowledge, Reinforcement Learning, Computational Modeling

Introduction

Syntactic Priming (SP, also known as "Structure Priming") is the linguistic phenomenon by which speakers tend to re-use syntactic structures across utterances (Bock,1996). Its existence is often touted as the strongest evidence that the same syntactic mechanisms are used in both language comprehension and language production. As such, manipulations that affect SP can be used to gather insight into how brain perceives, represents, and applies syntactic structures. For example, two notable studies (Loebell & Bock,2003;Hartsuiker, Pickering, & Veltkamp,2004) have show that that SP effects occur across languages, demonstrating that syntactic structure is represented in a way that is language-independent.

In this paper, we will use a novel manipulation of SP effects to investigate whether syntactic structures are represented within declarative or procedural memory. Our results, backed by computational models, strongly suggest that SP is based on procedural representations, and that these representation are learned and refined through Reinforcement Learning.

Background

In the past few decades, many researchers have attempted to determine the most likely mechanistic explanation for SP (Hartsuiker et al.,2004;Reitter, Keller, & Moore,2011;Chang, Dell, & Bock,2006). Experimental studies show that a range of factors could impact the strength of priming. For example, the priming effect is enhanced by the presentation of multiple primes, which is referred as the cumulativity of SP (Jaeger & Snider,2008). Not only the occurrence of primes matters, the lexical overlapping between prime and target also enhances priming, which is known as the lexical boosting effect (Pickering & Branigan,1998). Moreover, there is evidence for an inverse frequency interaction, showing that that the less frequently used syntactic structures are associated with with stronger priming effects (Jaeger & Snider,2008).

These effects have been used to support different underlying mechanisms that might account for SP. A main source of disagreement between these putative mechanisms is whether syntactic processing is relying on declarative or procedural representations. A group of researchers, for example, advocate a short-term residual activation mechanism account (Snider,2008;Jaeger & Snider,2008;Pickering & Branigan,1998) that implies a declarative representation, while another group of researchers believe that syntactic persistence is depending on implicit learning mechanisms (Chang et al.,2006;Bock & Griffin,2000) that point to a procedural representation. By incorporating both short-term activation account and long-term implicit account, a further dual mechanism account, Declarative/Procedural model of language is proposed by (Ullman,2004). Based on different mechanisms, different computational models have been developed to account for structural priming effects.

Most psycholinguistic studies have investigated syntactic priming effects using carefully controlled experimental items, ensuring that the linguistic stimuli have no mistakes and are produced flawlessly. However, in natural conversation, disfluencies and errors are very common when people are speaking. Usually, erroneous message is considered as interference that either slows down the processing or impedes peoples comprehension. Speech errors include ungrammatical construction, inappropriate word choice, ambiguous meaning, or absolute nonsense. Even though people may ignore minor speech errors in daily conversation, there is evidence that erroneous information does affect language processing, and might provide a further cue to the underlying representation of syntax. For example, people often change their mind and correct themselves mid-sentence while speaking. Slevc and Ferreira (2013) examined the priming effect in the context of correcting speech errors. They found that SP is significantly reduced when primes are corrected to the alternative syntactic structure.

The prediction error (i.e., surprise) associated with the syntactic structure of prime also affects subsequent language processing. There was evidence that the more surprising the prime is, which means higher prediction errors, the more likely to expect the same structure would occur later (Jaeger & Snider,2008).

The role played by errors in SP introduces a third point of view on the nature of SP, which can be catalogued under the "procedural" account. According to this point of view, syntactic structures are represented procedurally and their selection is guided by their perceived utility in terms of Reinforcement Learning, i.e., their estimated future amount of "rewards" or positive feedback signals (Sutton, Barto, et al.,1998). It is widely accepted that procedural knowledge, in general, is refined in a Reinforcement Learning-like manner through the backpropagation of reward or feedback signals. In fact, procedural knowledge and reward signals share the same computational substrate, in the dopamine-rich basal ganglia (Schultz, Dayan, & Montague,1997;Yin & Knowlton,2006). Furthermore, although the basal ganglia are not considered part of the cortical language network, an increasing number of studies have shown their involvement in language processing (Friederici,2006;Stocco, Yamasaki, Natalenko, & Prat,2014).

The connection between reward signals and procedural knowledge is apparent in some prominent general theories of cognition. For example, in the ACT-R cognitive architecture (Anderson,2009;Anderson et al.,2004), procedural knowledge is represented as production rules or simply *productions*, and productions are typically used to represent syntactic micro-operations in ACT-R models of language processing (Lewis & Vasishth,2005;Stocco & Crescentini,2005). But, in ACT-R, productions are selected on the basis of their expected , a scalar quantity that represents future rewards and is updated through repeated feedback signals according to a standard Reinforcement Learning rule:

$$U_{t+1}(p) = U_t(p) + \alpha \times (R_t - U_t(p)) \tag{1}$$

where $U_t(p)$ represents the utility U of production p at time point t.

In the case of linguistic phenomena, feedback signals could be provided directly by the process of successfully comprehending or producing a sentence. Thus, according to this view, SP would be the effect of increased utility of a syntactic structure following its successful use in comprehension.

If that is the case, we expect that ungrammatical sentences, in which rules are applied *unsuccessfully* and lead to a error signal and a re-analysis of a sentence, would result in negative feedback signals. These negative feedback signals would ultimately *decrease* the utility of the corresponding production, thus making the application of the same syntactic structure less likely to occur.

In this study, we set forward to test this alternative, RL-based account for syntactic priming, and to answer the question of whether perceiving incorrect linguistic information such as ungrammatical syntactic constructions would affect peoples subsequent language representation, particularly in syntactic choices of production. Furthermore, we will attempt to explain the observed patterns under Reinforcement Learning theory and simulate the behavioral results using ACT-R model.

Theoretical Hypotheses

Based on the proposed theories of SP, we can derive three different predictions about the effect of syntactically incorrect primes (See Figure 1). Across all predictions, we expect that syntactic priming effect will occur regardless of syntactic correctness. Specifically, the proportion of producing same construction is expected to be higher than producing alternative construction. We also expect that the priming effect will be different depending on whether the syntactic structure of prime is correct or not.

According to a purely declarative model (as exemplified, for instance, by Reitter's 2011 model), an ungrammatical prime should not have any differential SP effect than a grammatical one. In as much as the prime sentence can be correctly interpreted despite the syntactic error (and, in our experiment, we made sure this is the case), the same grammatical structure would be retrieved, thus causing the same activation boost for subsequent use. Thus, our Hypothesis 1, driven by the repetition between prime and target, states that there is no difference between grammatical and ungrammatical primes.

According to procedural, prediction-driven model (as exemplified by Jaeger & Snider,2008 and Snider,2008's exemplar-based model), the ungrammatical prime, being a low-frequency and unexpected structure, would generate greater surprisal and therefore *enhance* priming effect for same constructions production, but to weaken priming in alternative construction production. Specifically, Hypothesis 2 states that priming with ungrammatical sentence makes people more likely to produce same constructions, and less likely to produce alternative constructions than priming with grammatical one. Finally, according to our procedural/RL account, SP is due to the update of the perceived utility of a procedural syntactic structure, which is increased for successfully parsed (grammatical) sentences and decreased following unsuccessfully parsed (ungrammatical) ones. Driven by reward, Hypothesis 3 predicts an opposite pattern as Hypothesis 2, stating that, priming ungrammatical sentences is expected to increase the likelihood of producing alternative structures than those used in the priming sentences.

To explicitly formulate our hypothesis, we implemented it as an ACT-R model[1]. The model performs a simplified version of canonical SP task, first comprehending a sentence (in

[1]The code for all the models described in this paper is available on our laboratory's GitHub page: http://github.com/UWCCDL/SyntaxPriming

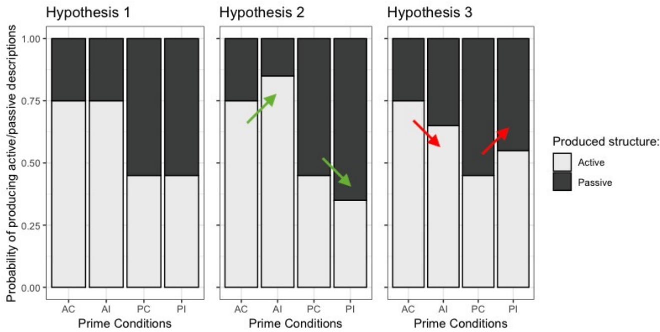

Figure 1: Three hypotheses driven by on different predictions. (white indicates active-form prime., gray indicates passive-form prime). Hypothesis 1: Declarative, driven by activation, predicts no effect of syntactic errors. Hypothesis 2: Driven by expectations, it predicts enhanced priming for (unexpected) ungrammatical sentences. Hypothesis 3: Driven by reward, it predicts reduced priming for ungrammatical sentences.

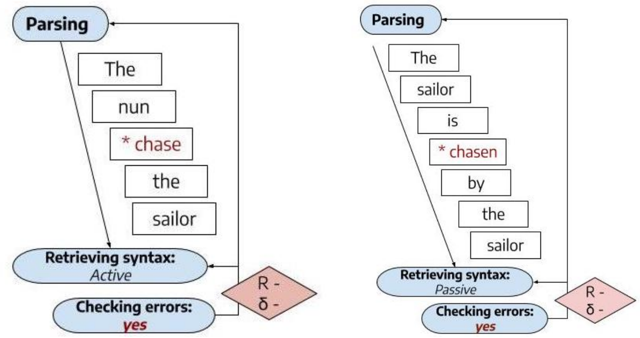

Figure 2: Two priming examples of the simple Reinforcement Learning model. Left: modeling AI priming. Right: modeling PI priming. White rectangles represents chucks encoding words. The blue rounded rectangles represents productions: *parsing* - parse in the prime; *retrieving syntax* - retrieve corresponding syntactic structure of the prime; *checking error* - check whether there is grammar errors in the prime. Diamond shapes represent feedback, either positive or negative. R indicates the reward term in Eq. 1, and δ reflects the reward prediction error term $R_t - U_t(p)$. When the model detects error, it sends a negative feedback signal to all the previous productions that have fired since the last reward. The predictions of this model are illustrated in Fig. 6

.

either active or passive form) and then producing a sentence to describe a picture. Both comprehension and production depend on the use of two production rules that implement the active and the passive sentence structures. In comprehension, these rules are used to mediate from the underlying sentence to its higher-level semantic representation. In language production, these rules are used to create a mental plan of the sequence of words to produce a description of the picture. Feedback signals are generated by detecting whether the comprehended sentence is grammatically correct or not. For simplicity, the process of parsing a sentence is drastically simplified (not unlike in Reitter et al.,2011), so that all the sentence information is available at once in a single visual "chunk" of information in ACT-R and feedback signals are only generated at the end of the comprehension process.

To examine the predictions of our model, we conducted a parameter space partitioning analysis of the model's behavior, and found that, across different initial utility values of the two syntactic structures and different reward values, the model produces the qualitative pattern of Figure 4.

To test between these alternative hypothesis, we conducted a novel SP experiment, introducing the novel manipulation of syntactic grammaticality of the priming sentences.

Materials and Methods

Participants

Ninety participants (35 female, 54 male, 1 other) were recruited online through Amazon Mechanical Turk, and performed the experiment in exchange for monetary compensation. Ethnicity includes 51.1% White, 36.7% Asian, 6.7% African American, 3.3% Latino or Hispanic American, and 2.2% Others. All participants were screened through a pre-experimental survey that gathered information about their language experience and background; only native English speakers without any history of brain damage, reading problems, nor language-related disorder were allowed to proceed to the experiment. Twenty-one were later excluded for failing to construct complete sentences in the language production task. The experimental protocol and inclusion criteria were approved by the Institutional Review Board at the University of Washington.

Materials

This picture description task is modified based on Hardy, Messenger, and Maylor's experiment (2017). A total of 36 trials with prime target pairs were created. Each picture is depicting a ditransitive action involving an agent and a patient. The verb of the action is printed under each picture. The prime sentence is either active-tense form grammatically correct (AC), passive-tense form grammatically correct (PC), active-tense form grammatically incorrect (AI) or passive-tense form grammatically incorrect (PI).

Ungrammatical prime sentences in the Passive Incorrect syntax condition (PI) were generated using seemingly correct but non-existing past participles modeled after existing verbs, such as "chasen" instead of "chased", "slapt" instead of "slapped", and "shooted" instead of "shot". In half of the trials within each condition, (N = 18 total), the prime picture and prime sentence are perfectly matched, while in the other half, the prime sentence is modified as semantically incorrect by which the identity of either agent or patient is wrong. This latter manipulation was designed to both make sure that participants were performing the task correctly and to separately measure the effect of syntactic errors from semantic errors.

Design

This study is a $2 \times 2 \times 2$ within-subject design, with three the factors being prime syntax (active vs. passive), grammatical correctness (correct vs. incorrect), and semantic correctness (correct vs. incorrect). In our notation, 4 syntax conditions: AC, AI, PC, PI x 2 semantic conditions: SC (semantically correct) and SI (semantically incorrect). Because, previous studies have demonstrated a stronger syntactic priming effect as prime and target are overlapping (Pickering & Branigan,1998), in this study prime and target always share the same action verb. The combination of three independent variable pairs are pseudo-randomized so in each syntax condition (AC, AI, PC, PI), each verb only occurs once, and each verb is modified as both semantic-correct and semantic-incorrect form.

Procedure

Most SP experiments make use of realistic, in-person dialogue between two participants, one of which is a confederate. The confederate verbally utters the primes and the participants responses are recorded for transcription. To simulate this seemingly realistic dialog situation online, the study described here used deception to convince participants that they were paired with another online "partner" and they were to take turns providing a description for a sentence and verifying the accuracy of their partner's description. In fact, there was no paired partner and all sentences typed by the partner were decided beforehand. At the end of the study, participants were fully debriefed about the use of deception.

In the online task, participants see a prime picture and are asked to verify whether the sentence constructed by the partner was correctly describing the picture or not. Followed by the verification task, there is a picture description task (see Figure 3). In the picture description phase, a picture and an appropriate verb are given, and participants need to type a sentence to describe the picture using the given verb. Participants are told that the game is proceeding in which the partner and the participant alternate between verifying if sentence-picture pair is matching, and constructing a sentence to describe the picture to the other. The game sets a randomly generated waiting time to simulate the amount of time needed by the fictional partner to type their own description.

The participant needs to complete a pre-screen survey that only eligible ones can continue. After giving consent, participants begin with a three-trial practice phase to familiarize themselves with the procedure. Between verification task and picture description task, the game sets a randomly generated waiting time to simulate verifying period of the "partner". At the end of the study, participants are given the debrief about the deception involved and are asked to complete a post-experiment survey.

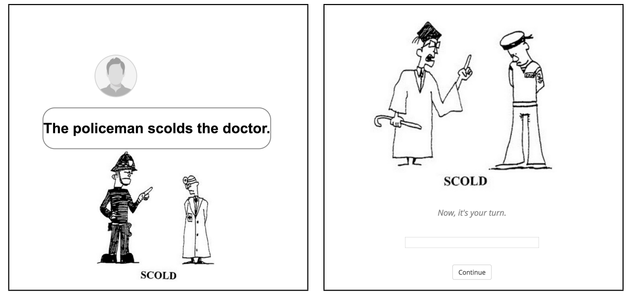

Figure 3: Example trial from the Online ASP Task. Left: During the Verification phase, subjects are asked to verify the congruence between the sentence and the picture. Right: During the Production phase, participants are asked to describe the target picture by typing a complete sentence that contains the given verb.

Results

The responses typed by participants were automatically analyzed with the Natural Language Toolkit (NLTK) package in Python and double-checked manually. Thirteen responses that could not be coded as neither active nor passive sentence were removed for data analysis.

The total of 2507 responses yield 75.79% active descriptions, 24.21% passive-voice. The analysis is conducted with the proportion of producing active out of active and passive responses. As expected, there is a significant effect of syntactic priming (Active vs. Passive), $F(1, 69) = 59.52$, $p < 0.001$, and a significant main effect of syntactic correctness (Syntax-Correct vs. Syntax-Incorrect), $F(1, 69) = 13.28$, $p = 0.001$. As expected, we find that there is no significant effect of semantic correctness on syntactic production. $F(1, 69) = 1.37$, $p = 0.25$.

Post-hoc analyses for significance indicate that the mean proportion of active descriptions is significantly lower in PC condition ($M = .64$, $SD = .32$) than that in PI ($M = .69$, $SD = .34$), $F(1, 69) = 5.05$, $p = 0.03$. The mean proportion of active descriptions is also significantly lower in the AC prime condition ($M = .84$, $SD = .24$) than in the AI conditions ($M = .8$, $SD = .21$), $F(1, 69) = 6.09$, $p = 0.01$.

As for the accuracy in verification task, overall accuracy rate is 79.92%. We find a significant effect of syntactic correctness on the accuracy rate $F(1, 68) = 57.66$, $p < 0.001$. People tend to verify picture more accurately when the sentence is grammatically correct ($M = .8796$, $SD = .16$) than the sentence is grammatically incorrect ($M = .72$, $SD = .21$). We also find that there is a significant effect of syntactic voice on the accuracy rate $F(1, 69) = 16$, $p = 0.001$. The accuracy of verification is significantly higher for active sentences ($M = .83$, $SD = .19$) than for passive sentences ($M = .77$, $SD = .21$). Interestingly, there is significant interaction effect on accuracy rate between syntactic correctness and syntactic voice, $F(1, 69) = 12.33$, $p = 0.001$.

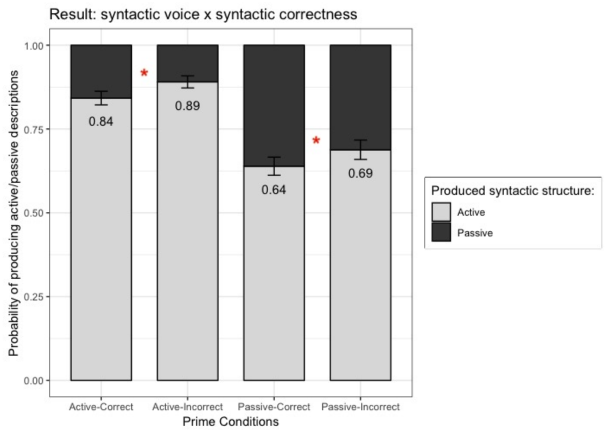

Figure 4: The proportion of active structures across conditions Asterisks "*" denote significant differences between conditions.

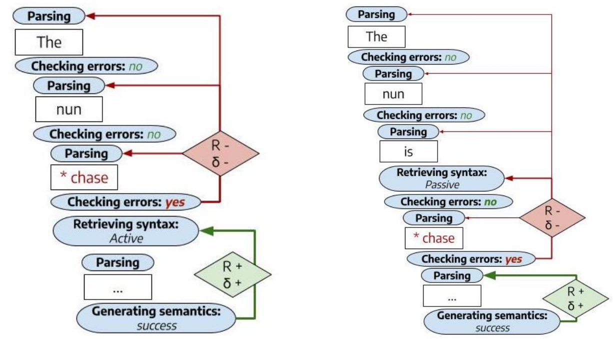

Figure 5: Two examples of the Sequential Procedural model. Left: Parsing of an ungrammatical, active sentence (AI). Right: parsing of an ungrammatical, passive (PI) sentence. This revised model explains both the activation boost found in AI priming and the activation drop in PI priming (see Fig. 6, Right)

Summary

Taken together, the results of our experiment provide a picture that is not entirely consistent with any of the previously discussed models, while the SP was present and robust (albeit less dramatic that in previous studies). Contrary to Reitter's model, there was a robust effect of syntactic grammaticality. These effects, however, did not comply precisely with either of the two competing accounts, that is, the procedural/expectancy and the procedural/RL hypotheses. In the passive sentences, an ungrammatical prime increased the likelihood of producing another active sentence, consistent. However, the data also show that semantic errors do not produce any effect, and, therefore, that the effect of errors can be localized to the processes of syntactic parsing.

A Sequential Procedural Model

One possible explanation for the lack of correspondence between the experimental results and our model is that the our procedural model was too naïve and did not appropriately take into account the different ways in which active and passive sentences are parsed. To explore this issue, we created a second computational model (See Figure 5).

This second model closely follows the structure of the procedural model described above. However, the new model simulates, at least partially, the sequential and incremental nature of sentence parsing. In particular, while the first model immediately detects the structure of the sentence (active vs. passive) and generates all feedback signals at the very end of the comprehension process, the second model delays the choice of the correct syntactic form until the first verb is encountered, and generates feedback signals both the end (when all sentences are successfully understood) and as soon as the first incorrect word is found (for ungrammatical ones).

This creates a novel asymmetry between the ungrammatical, active (AI) and ungrammatical, passive (PI) sentences. In the case of passive sentences, the first verb form encountered by the model is the word "is" (as in "the robber *is* chased (...)"); when the word "is" is encountered, the model can confidently select the production rule that encodes the passive structure. The grammatical mistake is then detected immediately thereafter (as in "the robber is *chasen* (...)"), thus generating an negative feedback that decreases the utility of the passive form. In this condition, therefore, the effect of grammaticality is identical to what was predicted by the previous model.

In the case of ungrammatical active sentences, however, the first verb form is also the first word for which a negative feedback signal can be generated (as in "the robber *chase* (...)"). In this case, the negative feedback is generated at the same time as the active sentence structure is selected, and, thus, does not affect the utility of the corresponding production. When the model successfully completes the sentence comprehension goal, a *positive* feedback signal is generated that propagates back to active form, thus increases its utility even if the sentence was ungrammatical.

This dynamic is further complicated by the fact that, to be selected in the face of a grammatically incorrect verb, a mechanisms of *procedural partial matching* had to be enabled. With this mechanism, productions are allowed to be selected even if their requirements are not perfectly satisfied. The price to pay for this imperfect selection is a temporary reduction in the associated utility. That is, instead of using a production's "true" utility $U_t(p)$, Eq. 1 uses the the reduced term $U_t^*(p)$:

$$U_t^*(p) = U_t(p) - MP(p) \tag{2}$$

where $MP(p)$ is the *mismatch penalty*, a fixed cost associated to applying a production rule to a condition in which not all the requirements are verified. This reduction reflects an intuitive greater uncertainty in the predicted future rewards for cases (such as ungrammatical sentences) in which productions are applied outside of their ideal conditions.

In turn, this reduced expectation affects the RL-based adjustments of utility. This is because these adjustments, according to the ACT-R theory and Eq. 1), reflect the magnitude of the reward prediction error δ_t, which is the difference between effective feedback signal and expected utility: $\delta_t = \alpha \times (R_t - U_t(p))$. It is easy to see that, for ungrammatical sentences, $U_t^* < U_t$ (because of the penalty match in Eq. 2) and, therefore, $\delta_t^* > \delta_t$: the as the utility U gets smaller, the adjustment δ_t gets larger, resulting in even greater benefit for the active form when it is selected while successfully parsing an ungrammmatical sentence.

To test our theory, we simulated the behavior of this model under different parameters. We found that the model consistently yields results consistent with our data. Fig. 6 depicts prototypical results (using $R = 1.0$, $\alpha = 0.2$, and $MP = 0.2$). Specifically, while the effect of grammaticality on the SP of passive sentences remains unchanged, the effect of grammaticality for active sentences either disappears (yielding equal probabilities of using the active form after a grammatical and an ungrammatical sentence) or results in *higher* rates of active sentences following ungrammatical primes.

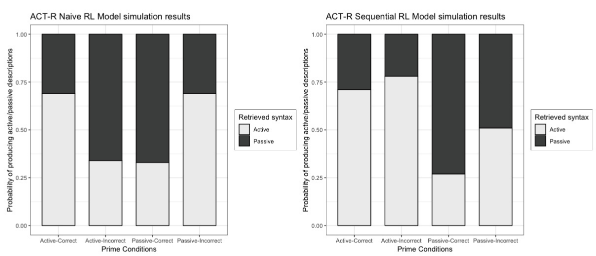

Figure 6: Simulation results from the two Reinforcement Learning models of syntactic priming. Left: Predictions of the simple, naïve models. Right: prediction of the Sequential model. The sequential model correctly predicts the general pattern of the experimental findings in Fig. 4.

Discussion

As demonstrated in many syntactic priming studies, people tend to re-use the same syntactic structures they are primed with. Consistent with this body of literature, our experiment shows a overall syntactic priming effect for active and passive structures, regardless of syntactic correctness and semantic correctness. This implies that the tendency of reproducing primed syntactic structures persists even if the linguistic information is noisy and erroneous.

In addition, our experimental results showed that syntactic priming is modulated by the grammaticality of the priming sentence. This result poses difficulty for purely declarative accounts (Reitter et al.,2011), which ascribe priming effects to the frequency and recency of syntactic structure retrieval. Furthermore, we found that this effect was specific to *grammatical errors*, and not to semantic errors (i.e., incorrectly labelled figures), thus restricting the effect to syntactic processes and excluding a general effect of surprise or attention to errors.

In general, these results support the idea that the priming effect of syntactic structures is dependent on procedural, rather than declarative memory, thus suggesting that syntactic structures are represented procedurally (Ullman,2004).

However, contrary to our expectations, our results show that participants actually always produce more active sentences following an ungrammatical sentence, regardless of the syntactic voice of the prime. This interesting pattern is at odds with both the declarative memory accounts (Hypothesis 1), the Activation Spreading account (Hypothesis 2), and a naïve procedural memory account (Hypothesis 3).

We found that this effect could be accounted for if our original naïve procedural model is expanded to included sequential parsing. Under these conditions, the order in which syntactic forms are selected and grammatical feedback signals are delivered becomes important. In particular, in ungrammatical active sentences, the negative feedback signal is delivered before the active form is selected, and the adjustment to the expected rewards of active structures is greater, thus reproducing the effects we found in our data.

Although successful, our second model is limited by the fact of having being designed post-hoc. To test its validity, the same experiment should be replicated using different syntactic structures, so that new predictions can be made and tested. Even within these limitations, however, the models describe herein have two important implications. First, they highlight the role of basic reinforcement learning mechanisms in learning, whose contribution might shed light on the basic computations underlying syntactic parsing as well as the contributions of subcortical structures to language (Hernandez et al.,2019). Second, our results highlight the importance of detailed computational models to explain psycholinguistic effects.

Acknowledgments

This research was made possible by a Top Scholar award from the University of Washington to YCY, and by a grant from the Office of Naval Research (ONRBAA13-003) to AS.

References

Anderson, J. R. (2009). *How can the human mind occur in the physical universe?* Oxford University Press.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological review*, *111*(4), 1036.

Bock, K. (1996). Language production: Methods and methodologies. *Psychonomic Bulletin & Review*, *3*(4), 395–421.

Bock, K., & Griffin, Z. M. (2000). The persistence of structural priming: Transient activation or implicit learning? *Journal of Experimental Psychology: General*, *129*(2), 177–192.

Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological Review*, *113*(2), 234–272.

Friederici, A. D. (2006). What's in control of language? *Nature neuroscience*, *9*(8), 991.

Hardy, S. M., Messenger, K., & Maylor, E. A. (2017). Aging and syntactic representations: Evidence of preserved syntactic priming and lexical boost. *Psychology and Aging*, *32*(6), 588–596.

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.

Hernandez, A. E., Claussenius-Kalman, H. L., Ronderos, J., Castilla-Earls, A. P., Sun, L., Weiss, S. D., et al. (2019). Neuroemergentism: A framework for studying cognition and the brain. *Journal of neurolinguistics*, *49*, 214–223.

Jaeger, T. F., & Snider, N. (2008). Implicit learning and syntactic persistence : Surprisal and cumulativity. *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, 827812).

Lewis, R. L., & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive science*, *29*(3), 375–419.

Loebell, H., & Bock, K. (2003). Structural priming across languages. *Linguistics*, *41*(5).

Pickering, M. J., & Branigan, H. P. (1998). The representation of verbs: Evidence from syntactic priming in language production. *Journal of Memory and language*, *39*(4), 633–651.

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

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599.

Slevc, L. R., & Ferreira, V. S. (2013). To err is human to structurally prime from errors is also human. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(3), 985–992.

Snider, N. (2008). An exemplar model of syntactic priming. unpublished doctoral dissertation, stanford university.

Stocco, A., & Crescentini, C. (2005). Syntactic comprehension in agrammatism: A computational model. *Brain and Language*, *95*(1), 127–128.

Stocco, A., Yamasaki, B., Natalenko, R., & Prat, C. S. (2014). Bilingual brain training: A neurobiological framework of how bilingual experience improves executive function. *International Journal of Bilingualism*, *18*(1), 67–92.

Sutton, R. S., Barto, A. G., et al. (1998). *Reinforcement learning: An introduction* (Vol. 135). MIT Press.

Ullman, M. T. (2004). Contributions of memory circuits to language: the declarative/procedural model. *Cognition*, *92*(1-2), 231–270.

Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience*, *7*(6), 464.