A neural network simulation of event-related potentials in response to syntactic violations in second-language learning

Stephan Verwijmeren (stephan.verwijmeren@ru.nl)
Department of AI, Radboud University
Nijmegen, The Netherlands

Stefan L. Frank (stefan.frank@ru.nl)
Centre for Language Studies, Radboud University
Nijmegen, The Netherlands

Hartmut Fitz (hartmut.fitz@mpi.nl)
Max Planck Institute, Radboud University
Nijmegen, The Netherlands

Yung Han Khoe (yunghan.khoe@ru.nl)
Centre for Language Studies, Radboud University
Nijmegen, The Netherlands

Abstract

Event-related potentials (ERPs) are used to study how language is processed in the brain, including differences between native (L1) and second-language (L2) comprehension. In low-proficiency L2 learners, syntactic violations give rise to an N400, but this changes into a P600 as their L2 proficiency increases. The precise functional interpretation of ERPs, however, remains a matter of debate. Fitz and Chang (2019) proposed a theory where ERPs reflect learning signals that arise from mismatches in predictive processing. These signals are propagated across the language system to make future predictions more accurate. We test if this theory can account for the N400-to-P600 switch in late bilinguals, by implementing a model capable of simulating the N400 and P600. We perform an experiment designed to elicit a P600 effect in simulated L2 learners progressing through learning stages. Simulated Spanish-English participants showed similar ERP effects in their L2 (English) as human participants did in ERP studies. Over the course of L2 learning, simulated N400 size decreased while P600 size increased, as it does in humans. Our findings support the viability of error propagation as an account of ERP effects, and specifically of how these can change over L2 learning.

Keywords: Event-related potential; N400; P600; prediction error; bilingualism.

Introduction

Psycholinguistic studies investigating neural mechanisms underlying adult second-language (L2) learning and processing often use electroencephalography (EEG), a technique for recording electrical voltage potentials produced by neural activity. Recorded potentials can be analyzed in relation to cognitive events, and can yield interpretable patterns called event-related potentials (ERPs) (Morgan-Short, 2014). ERP effects have been observed in response to syntactic violations in first language (L1) processing, as an increased positivity in the ERP waveform that starts around 600 ms after observing an anomalous word, as compared to its correct counterpart (Osterhout and Mobley, 1995). This effect is called a P600. Another ERP effect is reliably elicited in response to a lexico-semantic violation. This effect, called an N400, is a negative voltage deflection around 400 ms after an anomalous word, as compared to a semantically appropriate word (Kutas and Hillyard, 1980).

ERP research has been done to find out if L2 learners show similar ERP effects as native speakers for morpho-syntactic and lexico-semantic processing. Research has shown that L2 learners can show native-like ERP waveforms for L2 grammatical features that are present in their L1 as well as for features unique to their L2 (Morgan-Short, 2014). ERPs of L2 learners differing in proficiency indicate that some learners progress through stages of syntactic learning, suggesting that there is an intermediate stage of learning between no L2 grammatical knowledge and grammaticalization (McLaughlin et al., 2010). The observed ERP effects differ between studies. Some L2 learning studies that investigated syntactic processing found an N400 for learners with low proficiency and a P600 for learners with high proficiency, suggesting that L2 learners might rely more on lexical processing at early learning stages (Alemán Bañón et al., 2014; Antonicelli and Rastelli, 2022; Díaz et al., 2016; Esfandiari et al., 2021; Grey, 2022; Mickan and Lemhöfer, 2020; Nichols and Joanisse, 2019; Osterhout et al., 2008; Tanner et al., 2013, 2014). Other related studies found a similar effect for proficiency but ERPs were biphasic at low proficiency levels, resembling an N400 followed by a P600. With increasing proficiency, the amplitude of the N400 decreased and the P600 amplitude increased but ERP waveforms remained biphasic to a degree (Bian et al., 2021; Bowden et al., 2013; Caffarra et al., 2015; Esfandiari et al., 2020; Grey et al., 2018; McLaughlin et al., 2010; Morgan-Short et al., 2012; Morgan-Short, 2014; Pélissier et al., 2015). In the majority of studies, L2 proficiency was the most important factor determining ERP profiles (Antonicelli and Rastelli, 2022; Caffarra et al., 2015; McLaughlin et al., 2010; Morgan-Short, 2014).

Here we are interested in whether L2 learning stages reflect on the ERPs in simulated participants like in human participants. We do so by taking a monolingual computational cognitive model of sentence production that has been used to explain ERPs, and extending it to the bilingual case.

Computational models of ERP effects

Several connectionist cognitive models have been proposed to explain the N400 ERP effect in sentence comprehension (see Eddine et al., 2022, for a review). Some of these take the magnitude of change in neural activation as a predictor of the N400 (Rabovsky et al., 2018) while others take the network's prediction error to account for N400 size (Brouwer et al., 2017; Fitz and Chang, 2019; Frank et al., 2015).

While a number of models can potentially explain the N400, the models by Brouwer et al. (2017) and Fitz and Chang (2019) are in addition able to model the P600. Specifically, Fitz and Chang (2019) used Chang's (2002) Dual-path model to show that prediction error corresponds to N400 size and backpropagated error corresponds to P600 size across a wide range of studies, providing support for the hypothesis that ERPs might reflect learning signals. This account of the N400 and P600 is known as the Error Propagation account.

The Dual-path model is a connectionist model of sentence production and syntactic development. The model has two pathways. The first pathway is the sequencing system that learns how words are ordered in a sentence and is based on the Simple Recurrent Network (Elman, 1990). The second pathway is a meaning system that learns how to map messages onto sentences in a target language. Previously, the Dual-path model was used to explain a wide range of sentence production phenomena in a number of different languages (Chang et al., 2006, 2015; Janciauskas and Chang, 2018; Tsoukala et al., 2017, 2021). For our studies, we used a bilingual extension of the Dual-path model (Tsoukala et al., 2021).[1]

The present study

We perform a computational modelling experiment to investigate whether simulated L2 learners progress through stages of syntactic learning, and further test the viability of Error Propagation as an account of ERPs. We do this by ascertaining whether a P600 effect can be simulated by the Bilingual Dual-path model, and whether the magnitude of this effect increases in later L2 learning stages. We simulate native speakers of Spanish (L1) who start learning English (L2) from a later age. At every L2 learning stage, we run a subject-verb number agreement experiment similar to one of the experiments in Fitz and Chang (2019), presenting simulated participants with stimuli containing syntactic violations that elicit a P600 in native speakers (Osterhout et al., 2008; Tanner et al., 2013, 2014), and with control sentences without such violations.

We expect to find a simulated P600 effect in the Bilingual Dual-path model, since Fitz and Chang (2019) were able to have the monolingual Dual-path model reproduce N400 and P600 effects for stimuli used in a number of human EEG studies. We further expect N400 and P600 effects to occur and their magnitude to decrease and increase, respectively, through learning stages, because ERP effects and their magnitude in L2 learners have been shown to be primarily determined by proficiency (Antonicelli and Rastelli, 2022; Caffarra et al., 2015; McLaughlin et al., 2010; Morgan-Short, 2014). We specifically expect the P600 effect to be more pronounced at later learning stages since advanced L2 learners show native-like ERP waveforms for L2 grammatical features (Morgan-Short, 2014). Additionally, we specifically expect the N400 effect to decrease in magnitude at later learning stages, because lexical learning precedes syntactic learning in L2 learners and L2 learners seem to rely on lexical processing early on because of this (McLaughlin et al., 2010).

Methods

To simulate late Spanish-English bilinguals, we trained the Bilingual Dual-path model (Figure 1) to learn Spanish from "infancy" and English as L2 at a later stage. The training input to the model consisted of sentences from two artificial languages (modelled on Spanish and English) that were paired with messages that encoded their meaning. The model learned to express messages as sentences of the target language (Spanish or English) by predicting the next word.

Artificial languages Table 1 shows the different constructions in the artificial languages. Constructions were distributed uniformly in the training input. Taken together, the two artificial languages consisted of 258 lexical items: 121 nouns, 11 adjectives, 6 pronouns, 6 determiners, 12 prepositions, 87 verbs, 7 auxiliary verbs, 6 verb inflectional morphemes, 1 plural noun marker, and the period. The inflectional morphemes were used to generate verbs with simple, progressive and perfect aspect in present or past tense. The plural noun marker was used to generate plural nouns.

The meaning space had 116 concepts and 7 thematic roles. Thematic roles are similar to those from Chang et al. (2006). To provide a simple example, the meaning of "the old lady carves a cake" would be represented as AGENT: LADY; ACTION-LINKING: CARVE; PATIENT: CAKE; AGENT-MODIFIER: OLD. This is implemented by introducing fixed-weight connections between role units and concept units (see Figure 1).

Model configuration and training For our simulations, we modified the Bilingual Dual-path model to resemble the architecture used in Fitz and Chang (2019): Previous word-history and role-history layers were added to the model which kept a running average of the activation of the input layer and role layer, respectively, and were connected to the hidden layer.

As pre-registered[2], all models used 50 hidden-layer units and 30 compress-layer units. Internal layer units used the logistic activation function; the output layer units used a softmax activation function. Weights were initialized randomly, uniformly between ± 1. Fixed weights for concept-to-role

[1] https://gitlab.com/yhkhoe/bilingual-dual-path/-/tree/ICCM2023

[2] The pre-registration can be accessed here: `https://aspredicted.org/blind.php?x=CGL_X3R`

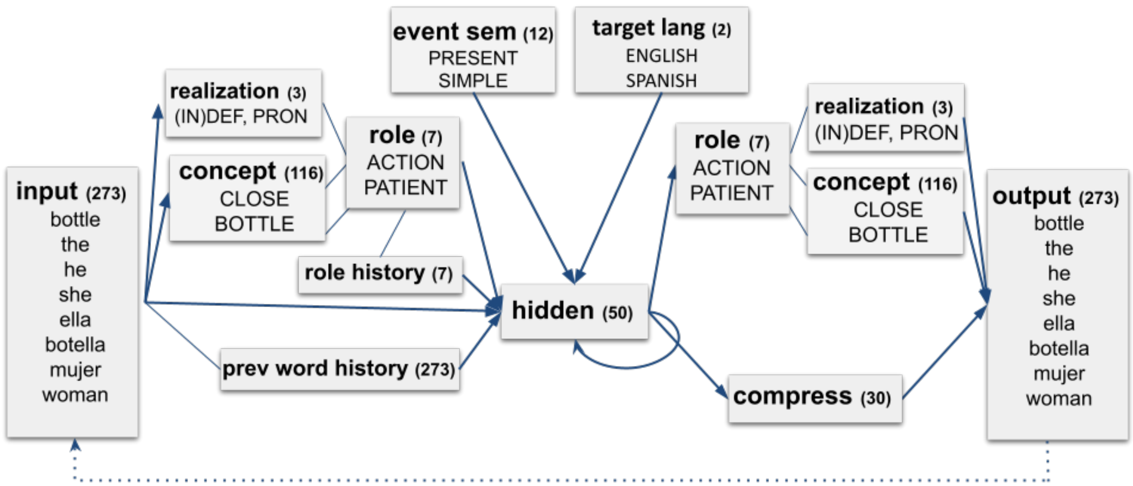

Figure 1: Architecture of the Bilingual Dual-path model. The model learns to map messages onto sentences in different languages by predicting the next word in its input. The sequencing system (lower path) maps from the input through a hidden layer to the output via a compression layer. The meaning system (upper path) uses information about thematic roles, concepts, and the realization of concept (e.g., by a pronoun or with an (in)definite determiner). The number of units per layer are shown in parentheses. Figure adapted from Tsoukala et al. (2021).

connections and realization-to-role connections were set to a value of 6. The concept layer had a set bias of -3.

As pre-registered, for each of 60 model subjects and for Spanish and English combined, we generated 10,000 unique message-sentence pairs for training and a novel set of 200 message-sentence pairs for testing. The sentences are approximately equally divided over the two languages, where the percentage of Spanish sentences was sampled from a uniform distribution between 48% and 52% and the rest was English. Following Fitz and Chang (2019), the message was excluded from 70% of the training items. Each model first iterated five times over its monolingual Spanish training set, followed by 75 epochs over its bilingual training set. The training set's order was randomized at the beginning of each of these 80 epochs. The model learned by steepest descent backpropagation, with momentum set to 0.9. Initially, the learning rate was set to 0.1, it decreased linearly to 0.02 over the 5 epochs of monolingual training, and then stayed constant during bilingual training.

Model evaluation After each epoch, model accuracy was tested using a 200-sentence test set. The model's L2 English proficiency was evaluated with two measures. First, syntactic accuracy was measured as the percentage of sentences for which all words had the correct part of speech. Second, meaning accuracy was measured as the percentage of syntactically correct sentences that also conveyed the target message without additions. As pre-registered, we excluded the 20 subjects with the lowest meaning accuracy, leaving data from 40 model subjects.

Experimental trials To elicit ERPs, we generated 30 English sentence pairs, each consisting of a control and a violation item. The control was an active transitive sentence where the verb form agreed with the subject in number. In the violation item, the verb did not agree with subject number. Violations were created by adding or omitting the inflectional marker for singular verbs (-ss), see Table 2.

Model subject differences Weights are initialized randomly, and differed between subjects. The percentage of Spanish versus English (training and testing) sentences varied between subjects, ranging from 48/52 to 52/48. The distribution of constructions is the same for all subjects. Training, testing and experimental trial sentences in the same language with the same constructions can differ between subjects in two ways. Firstly, sentences can differ in content-words resulting in different meaning of sentences. Secondly, singular nouns can differ in definiteness of the article.

Measuring model ERPs After every training epoch, the model was tested on the experimental sentence pairs. As in Fitz and Chang (2019), learning was turned on in the model during processing, but connection weights were reset to the weights of the respective training epoch after each test sentence in order to exclude learning effects during the experiment. The state in which the model encountered each trial was thus the same for all of the sentences.

We measured the prediction error at the output layer and the hidden layer (see Fitz and Chang, 2019, for details). The prediction error of output unit j is the difference between its activation y_j and the target activation t_j, or: $\delta_j = y_j - t_j$, with $y_j \in [0,1]$ and $t_j \in \{0,1\}$. This error was backpropagated in the network, as happens during training, to generate error at deeper layers. Error for units connected to the output layer was calculated as shown in Eq. 1, where k indexes the units connected to the output layer with weight w_{kj}, and j references

Table 1: Constructions with English example sentences. In the artificial language modelled on English, inflectional morphemes -prg, -prf and -ss are used for verb conjugations in progressive, perfect, and 3rd-person present simple tense, respectively.

Construction	Example sentence
Animate intransitive	The woman is play -prg
Animate with intransitive	The woman is play -prg with a dog
Inanimate intransitive	The apple is fall -prg
Locative	The boy is walk -prg around the school
Theme-experiencer (active)	The uncle surprise -ss the grandfather
Theme-experiencer (passive)	The grandfather is surprise -prf by the uncle
Transitive (active)	The girl bake -ss a cake
Transitive (passive)	The cake is bake -prf by the girl
Cause-motion	The hostess is put -prg a cactus into the office
Benefactive transitive	The grandmother repair -ss the cup for the girl
State-change	The waiter is fill -prg the cup with water
Locative alternation	The man spray -ss the sink with water

the units that are backpropagating error.

$$\delta_k = y_k(1-y_k)\sum_{j=1}^{n}\delta_j w_{kj} \quad y_k \in [0,1] \tag{1}$$

Error was calculated the same for other layers backpropagating error into the network. The error was collected after the transitive verb where the third-person singular morpheme was present or absent. The simulated N400 and P600 sizes are the sums over $|\delta|$ of the output- and hidden-layer units, respectively. Note that the scales of these two measures are not comparable because the output units, unlike the hidden units, use the softmax activation function and therefore their activations always sum to 1.

Table 2: Example sentences for the experimental trials. The bold morphemes indicate the sentence position where prediction error was measured.

Example sentence	Subject Nr	Agreement
the old lady carve **-ss** a cake	Singular	Control
the old lady carve **a** cake	Singular	Violation
the old lady -s carve **a** cake	Plural	Control
the old lady -s carve **-ss** a cake	Plural	Violation

Results

Figure 2 displays the proficiency of the model at the start and the end of bilingual training.

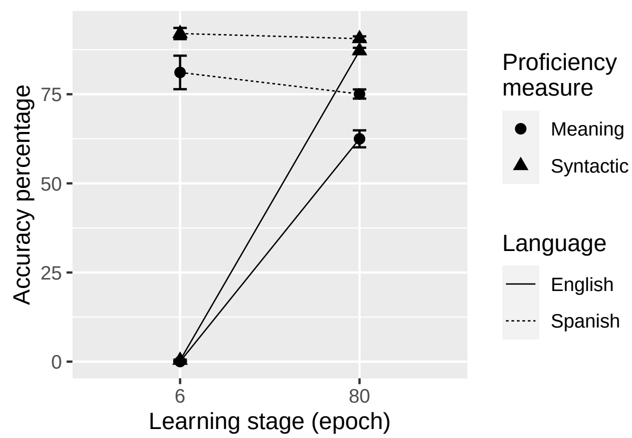

Figure 2: Mean proficiency of model. The syntactic and meaning accuracy are displayed for the first and last epoch of bilingual training. The error bars show the 95% confidence interval.

The mean prediction error over L2 learning stages at the hidden layer and the output layer are displayed in Figure 3, respectively. At the output layer, the mean error (simulating N400) for the VIOLATION items, was 1.89 at the start of bilingual training and increased to 1.93 at epoch 19, whereafter it decreased to 1.33 over the learning epochs. The mean error at the hidden layer (simulating P600) for the VIOLATION condition was 3.30 at the start of bilingual training, and increased over the learning epochs to 12.52. For the CONTROL items, error at both layers was high initially, but decreased to values close to 0 during L2 learning.

Pre-registered analysis

As pre-registered, we analyzed the data from our experiment with a linear mixed-effects model, using the lmer function from the package lme4 (Bates et al., 2015) in R (R Core Team, 2013). The model fits the prediction error from the Bilingual Dual-path model, a numerical value. The regression model[3] included the predictors of interest: AGREEMENT, LAYER, LEARNING_STAGE and their interactions. AGREEMENT and LAYER were sum-coded. AGREEMENT levels Control and Violation were coded -1 and $+1$, respectively. Levels Hidden and Output of LAYER were coded $+1$ and -1, respectively. The number of L2 training epochs is indicated by the LEARNING_STAGE predictor, which was standardized. We fit random intercepts for model participants, and by-participant random slopes for the three predictors of interest and their interactions. Table 3 reports estimates, 95% confidence intervals,

[3]The script for the mixed-effects model can be accessed here: `https://osf.io/yprjk/?view_only=aae2b8a52819475eb127721931de19ba`

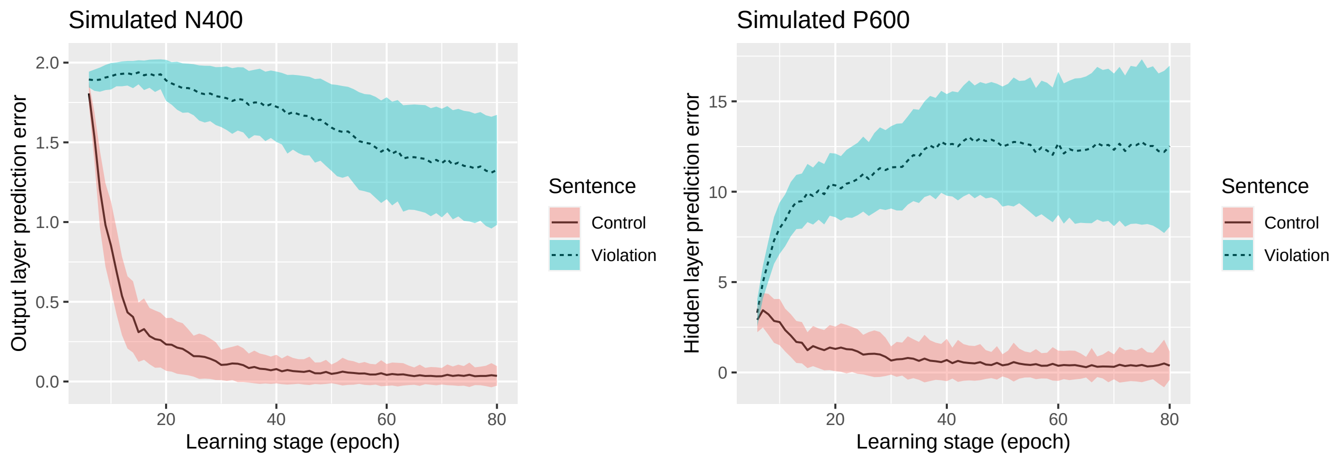

Figure 3: Mean prediction error (averaged over all model subjects) as a function of learning stage, in the output layer (left panel) and in the hidden layer (right panel), for number agreement violation and control items,. Shaded areas represent the 95% CI.

Table 3: Summary of the fixed effects in the linear mixed-effects models.

Predictor	Est.	95% CI	SE	df	t-value	$\Pr(>\lvert t \rvert)$
Intercept	3.54	$[3.32, 3.75]$	0.11	40.00	33.84	<0.001
AGREEMENT	3.00	$[2.81, 3.20]$	0.10	40.05	30.76	<0.001
LAYER	2.61	$[2.41, 2.82]$	0.10	40.00	26.17	<0.001
LEARNING_STAGE	0.10	$[-0.04, 0.24]$	0.07	40.17	1.41	0.165
AGREEMENT:LAYER	2.28	$[2.10, 2.46]$	0.09	40.04	25.49	<0.001
AGREEMENT:LEARNING_STAGE	0.50	$[0.36, 0.63]$	0.07	40.15	7.31	<0.001
LAYER: LEARNING_STAGE	0.31	$[0.19, 0.43]$	0.06	40.18	5.08	<0.001
AGREEMENT:LAYER: LEARNING_STAGE	0.49	$[0.37, 0.61]$	0.06	40.16	8.34	<0.001

standard errors, degrees of freedom, t-values and p-values.

The positive estimate for the interaction between the predictors AGREEMENT, LAYER and LEARNING_STAGE (Estimate = 0.49, 95% CI = [0.37, 0.61]) indicates that the learning stages affect the two layers' sensitivity to violated sentences differently. The estimate has a confidence interval not including zero, thus there was an effect of the three-way interaction between these predictors. As Figure 3 clearly shows, this interaction is driven by an increasing effect of violation in the hidden layer combined with a decreasing effect of violation in the output layer.

Discussion

In the present work, we investigated whether simulated L2 learners progress through stages of syntactic learning. We used a connectionist model of syntactic development (Chang, 2002) to simulate Spanish-English bilinguals and exposed the model to L2 number-agreement violations at different points in time. Similar to the account in Fitz and Chang (2019), we recorded ERPs in response to these syntactically anomalous sentences from the model. On this account, ERPs are summary signals of brain activity that index the propagation of prediction error during comprehension whose functional role is to support learning. Prediction error at the output layer was used to model the N400 and the backpropagated prediction error at the hidden layer was used to model the P600. The results of our simulations revealed a clear P600 effect for syntactically anomalous sentences in the L2, as well as a clear N400 effect early in acquisition. We also found that over time the P600 increased as the model became more proficient in the L2 and the N400 decreased over time. These findings are similar to human L2 learners as reported in several ERP studies on second language acquisition (Antonicelli and Rastelli, 2022; Caffarra et al., 2015; McLaughlin et al., 2010; Morgan-Short, 2014). Thus, our results support a theory of stages of syntactic learning in L2 learners where the magnitude of different ERP components changes during acquisition.

In our simulations, monolingual training resulted in optimal network weights for the L1, after which new L2 learning required a considerable amount of further training. At the beginning of L2 learning, the model does not know the English syntax for noun-verb number agreement. Consequently, after seeing the verb, the model activates a variety of candidate words and morphemes, which leads to large prediction error at the lexical output layer, and thus a large-amplitude N400 prediction for both violations and control sentences.

Prediction error at the hidden layer indexing the P600, in contrast, is relatively small because the model has not yet learned the syntax of agreement. As the model gradually acquires agreement, word predictions after the verb become increasingly more accurate because they are more and more driven by learned syntactic knowledge in the hidden layer. When the model is presented with a number agreement violation item, there is now a larger mismatch between the observed violation and the correct word predictions made by the model at this sentence position. Because the correct prediction is due to syntactic knowledge at the hidden layer, the hidden layer gets the majority of the blame when such a mismatch occurs. Thus, the size of the P600 effect increases during syntactic learning. The lexical output layer, on the other hand, gradually receives less blame as the syntax of agreement is acquired deeper in the network, which leads to a decrease in the N400 effect over time.

The error propagation account explains why ERPs elicited by lexical violations (N400) precede ERPs in response to syntactic violations (P600) and this account has been able to reproduce key findings from a considerable number of monolingual ERP studies (Fitz and Chang, 2019). The results presented here on bilingual ERPs, and how they change over development, adds further support for this account. Apart from the error propagation account, the model of Brouwer et al. (2017) can also explain monolingual N400 and P600 effects but it remains to be tested whether this model would be able to simulate ERP effects in bilinguals and the change in size of these effects during second language acquisition. What is unique about the error propagation account is that it can naturally model and explain ERPs in development because on this account ERPs are directly linked to learning. Therefore, the magnitude of ERP effects is expected to change as different pieces of linguistic knowledge are acquired. One limitation of the model is that it currently does not account for differences in the precise onset of the N400 or P600 and that it does not model earlier ERP components such as the early left-anterior negativity (eLAN) which has been elicited in some bilingual studies (Caffarra et al., 2015).

At present, it is unclear to what extent L1–L2 language similarity affects ERP effects in bilinguals. Some studies showed reduced P600 effects, or no P600 effect, for syntactic features that are instantiated differently between languages (Antonicelli and Rastelli, 2022; Liu et al., 2017; Morgan-Short, 2014), while other studies have shown P600 effects for syntactic L2 features regardless of L1–L2 similarity (Caffarra et al., 2015; McLaughlin et al., 2010; Morgan-Short, 2014). In future work, the proposed model will be used to shed more light on the role of language similarity in simulated bilinguals.

References

Alemán Bañón, J., Fiorentino, R., and Gabriele, A. (2014). Morphosyntactic processing in advanced second language (L2) learners: An event-related potential investigation of the effects of L1–L2 similarity and structural distance. *Second Language Research*, 30(3):275–306.

Antonicelli, G. and Rastelli, S. (2022). Event-related potentials in the study of L2 sentence processing: A scoping review of the decade 2010-2020. *Language Acquisition*, pages 1–38.

Bates, D., Mächler, M., Bolker, B., Walker, S., et al. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(i01).

Bian, J., Zhang, H., and Sun, C. (2021). An ERP study on attraction effects in advanced L2 learners. *Frontiers in Psychology*, 12.

Bowden, H. W., Steinhauer, K., Sanz, C., and Ullman, M. T. (2013). Native-like brain processing of syntax can be attained by university foreign language learners. *Neuropsychologia*, 51(13):2492–2511.

Brouwer, H., Crocker, M. W., Venhuizen, N. J., and Hoeks, J. C. (2017). A neurocomputational model of the N400 and the P600 in language processing. *Cognitive Science*, 41:1318–1352.

Caffarra, S., Molinaro, N., Davidson, D., and Carreiras, M. (2015). Second language syntactic processing revealed through event-related potentials: An empirical review. *Neuroscience & Biobehavioral Reviews*, 51:31–47.

Chang, F. (2002). Symbolically speaking: A connectionist model of sentence production. *Cognitive Science*, 26(5):609–651.

Chang, F., Baumann, M., Pappert, S., and Fitz, H. (2015). Do lemmas speak German? a verb position effect in German structural priming. *Cognitive Science*, 39(5):1113–1130.

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

Díaz, B., Erdocia, K., De Menezes, R. F., Mueller, J. L., Sebastián-Gallés, N., and Laka, I. (2016). Electrophysiological correlates of second-language syntactic processes are related to native and second language distance regardless of age of acquisition. *Frontiers in Psychology*, 7:133.

Eddine, S. N., Brothers, T., and Kuperberg, G. R. (2022). The N400 in silico: A review of computational models. *The Psychology of Learning and Motivation*, page 123.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14(2):179–211.

Esfandiari, L., Nilipour, R., Maftoon, P., and Nejati, V. (2021). Native-like event-related potentials in processing the second language syntax: Late bilinguals. *Caspian Journal of Neurological Sciences*, 7(2):51–59.

Esfandiari, L., Nilipour, R., Nejati, V., Maftoon, P., and Khosrowabadi, R. (2020). An event-related potential study of second language semantic and syntactic processing: Evidence from the declarative/procedural model. *Basic and Clinical Neuroscience*, 11(6):841.

Fitz, H. and Chang, F. (2019). Language ERPs reflect learning through prediction error propagation. *Cognitive Psychology*, 111:15–52.

Frank, S. L., Otten, L. J., Galli, G., and Vigliocco, G. (2015). The ERP response to the amount of information conveyed by words in sentences. *Brain and Language*, 140:1–11.

Grey, S. (2022). Variability in native and nonnative language: An ERP study of semantic and grammar processing. *Studies in Second Language Acquisition*, pages 1–30.

Grey, S., Sanz, C., Morgan-Short, K., and Ullman, M. T. (2018). Bilingual and monolingual adults learning an additional language: ERPs reveal differences in syntactic processing. *Bilingualism: Language and Cognition*, 21(5):970–994.

Janciauskas, M. and Chang, F. (2018). Input and age-dependent variation in second language learning: A connectionist account. *Cognitive Science*, 42:519–554.

Kutas, M. and Hillyard, S. A. (1980). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science*, 207(4427):203–205.

Liu, H., Dunlap, S., Tang, Y., Lu, Y., and Chen, B. (2017). The modulatory role of L1 and L2 morphosyntactic similarity during production of L2 inflected words: An ERP study. *Journal of Neurolinguistics*, 42:109–123.

McLaughlin, J., Tanner, D., Pitkänen, I., Frenck-Mestre, C., Inoue, K., Valentine, G., and Osterhout, L. (2010). Brain potentials reveal discrete stages of L2 grammatical learning. *Language Learning*, 60:123–150.

Mickan, A. and Lemhöfer, K. (2020). Tracking syntactic conflict between languages over the course of L2 acquisition: A cross-sectional event-related potential study. *Journal of Cognitive Neuroscience*, 32(5):822–846.

Morgan-Short, K. (2014). Electrophysiological approaches to understanding second language acquisition: A field reaching its potential. *Annual Review of Applied Linguistics*, 34:15–36.

Morgan-Short, K., Steinhauer, K., Sanz, C., and Ullman, M. T. (2012). Explicit and implicit second language training differentially affect the achievement of native-like brain activation patterns. *Journal of Cognitive Neuroscience*, 24(4):933–947.

Nichols, E. S. and Joanisse, M. F. (2019). Individual differences predict erp signatures of second language learning of novel grammatical rules. *Bilingualism: Language and Cognition*, 22(1):78–92.

Osterhout, L. and Mobley, L. A. (1995). Event-related brain potentials elicited by failure to agree. *Journal of Memory and Language*, 34(6):739–773.

Osterhout, L., Poliakov, A., Inoue, K., McLaughlin, J., Valentine, G., Pitkanen, I., Frenck-Mestre, C., and Hirschensohn, J. (2008). Second-language learning and changes in the brain. *Journal of Neurolinguistics*, 21(6):509–521.

Pélissier, M., Krzonowski, J., and Ferragne, E. (2015). Effect of proficiency on subject-verb agreement processing in french learners of english: An erp study. In *Proceedings of the International Conference of Experimental Linguistics, EXLing*.

R Core Team (2013). Core R: A language and environment for statistical computing. *R Foundation for Statistical Computing, Vienna, Austria. Online: http://www.R-project.org*, 201.

Rabovsky, M., Hansen, S. S., and McClelland, J. L. (2018). Modelling the N400 brain potential as change in a probabilistic representation of meaning. *Nature Human Behaviour*, 2(9):693–705.

Tanner, D., Inoue, K., and Osterhout, L. (2014). Brain-based individual differences in online L2 grammatical comprehension. *Bilingualism: Language and Cognition*, 17(2):277–293.

Tanner, D., McLaughlin, J., Herschensohn, J., and Osterhout, L. (2013). Individual differences reveal stages of L2 grammatical acquisition: ERP evidence. *Bilingualism: Language and Cognition*, 16(2):367–382.

Tsoukala, C., Broersma, M., Van den Bosch, A., and Frank, S. L. (2021). Simulating code-switching using a neural network model of bilingual sentence production. *Computational Brain & Behavior*, 4(1):87–100.

Tsoukala, C., Frank, S. L., and Broersma, M. (2017). "He's pregnant": Simulating the confusing case of gender pronoun errors in L2 English. In *the 39th Annual Meeting of the Cognitive Science Society*, pages 3392–3397. Cognitive Science Society.