UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

He's pregnant": simulating the confusing case of gender pronoun errors in L2English

Permalink

https://escholarship.org/uc/item/9sf3q33f

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 39(0)

Authors

Tsoukala, Chara
Frank, Stefan L.
Broersma, Mirjam

Publication Date

2017

Peer reviewed

"He's pregnant": simulating the confusing case of gender pronoun errors in L2 English

Chara Tsoukala (c.tsoukala@let.ru.nl)
Centre for Language Studies
Radboud University

Stefan L. Frank (s.frank@let.ru.nl)
Centre for Language Studies
Radboud University

Mirjam Broersma (m.broersma@let.ru.nl)
Centre for Language Studies
Radboud University

Abstract

Even advanced Spanish speakers of second language English tend to confuse the pronouns 'he' and 'she', often without even noticing their mistake (Lahoz, 1991). A study by Antón-Méndez (2010) has indicated that a possible reason for this error is the fact that Spanish is a pro-drop language. In order to test this hypothesis, we used an extension of Dual-path (Chang, 2002), a computational cognitive model of sentence production, to simulate two models of bilingual speech production of second language English. One model had Spanish (ES) as a native language, whereas the other learned a Spanish-like language that used the pronoun at all times (non-pro-drop Spanish, NPD_ES). When tested on L2 English sentences, the bilingual pro-drop Spanish model produced significantly more gender pronoun errors, confirming that pronoun dropping could indeed be responsible for the gender confusion in natural language use as well.

Keywords: L2 pronoun errors, language transfer, Dual-path model, bilingual sentence production

Introduction

Second language (L2) speech errors have been employed in the past as a means to understand bilingual speech production as well as the acquisition process of a foreign language (Antón-Méndez, 2010; Poulisse, 1999). Certain L2 errors are observed more often due to discrepancies between the first language (L1) and the L2. For example, if the expression of a message in the L2 requires the inclusion of a specific feature that would not be necessary in the L1, then speakers of these two languages may produce a speech error in their L2 due to L1 transfer (Odlin, 1989). In this study, we focus on a gender-related L2 pronoun error that has been observed among native speakers of Spanish and Italian; namely, errors involving the third person singular nominative pronouns 'he' and 'she'. Even advanced Spanish speakers of L2 English occasionally confuse the two pronouns, referring to an actress as 'he' or a father as 'she', often without even noticing their mistake (Lahoz, 1991). At first, this phenomenon seems surprising because the Spanish language does have two equivalent pronouns ('él' for 'he' and 'ella' for 'she'), and also a very strong separation between the two genders, even more so than in English. For instance, depending on the suffix a word can be feminine or masculine (e.g., maestro - teacher

[masculine], maestra - teacher [feminine]; niño - child [masculine, a.k.a. boy], niña - child [feminine, a.k.a. girl]). This means that the gender mistakes that Spanish speakers make in English cannot be attributed to the lack of familiarity with the distinction. Furthermore, the challenge the English pronoun system poses for native speakers of Spanish could not be due to its inherent difficulty, as this would mean that most non-native speakers of English, regardless of their L1, would produce the same mistake. As Lahoz (1991) noted, a low proficiency level of the native Spanish speakers is not a reason either. This was also demonstrated in the experiments of Antón-Méndez (2010), where the participants showed an intermediate to upper intermediate knowledge of English. Finally, note that the gender mismatch error cannot be classified as a syntactic error; the produced sentence is grammatically correct, but it conveys the wrong meaning.

A hypothesis which has been put forward (Lahoz, 1991; Antón-Méndez, 2010) regarding the cause of errors in the use of English pronouns is the pro-drop status of the Spanish but not the English language. In pro-drop languages, nominative personal pronouns are often omitted (1b) because the number and person information is conveyed in the conjugated verb (Davidson, 1996), whereas in English the omission of the pronoun would result in an ungrammatical sentence (2b).

1.(a) *Él/ella* tiene un perro (Spanish)

(b) tiene un perro

2.(a) *He/she* has a dog (English)

(b) * has a dog

It is hard to imagine, however, how the pro-drop feature of the L1 might result in a gender pronoun error ("He's walking", when referring to a woman) instead of an omission ("Is walking"), which would be the case in a direct language transfer.

As a matter of fact, native speakers of Spanish have been noted to produce another gender-related pronoun error in English, this time regarding possessive pronouns ('his', 'her'). Due to the high frequency of this type of errors, a lot more

emphasis has been given to the misuse of these pronouns than the subject pronouns (White, Muñoz, & Collins, 2007; Anton-Mendez, 2011). The reason that English possessive pronouns pose a challenge for native speakers of Spanish is most likely that in Romance languages the possessive pronoun agrees in gender and number with the possessum, namely the noun that follows, whereas in Germanic languages such as English the possessive pronoun refers to the antecedent. For example:

i. His daughters are on vacation.

   [*his*: 3rd person masculine singular]

ii. Sus hijas están de vacaciones.

   [*sus*: 3rd person feminine plural]

Due to the different information encoding Spanish speakers of English may occasionally make gender mistakes such as "He called her mother", where 'her' refers to the antecedent ('he') and not a different female person. This is because 'mother' is female, and a Spanish speaker would use that gender information to construct the possessive pronoun in Spanish. The resulting error in English is, of course, confusing, as a speaker of English would not guess that 'her' in this case refers to the same subject ('he'). The gender error in the case of L2 English possessive pronouns seems clearly due to L1 transfer, because the properties of Spanish are directly applied to English. In the case of the subject pronoun gender errors, on the other hand, it is not evident that the pro-drop feature of one language would lead to a gender error in L2. The present study addresses only the latter type of errors.

Antón-Méndez (2010) has investigated the hypothesis that the pro-drop feature of Spanish is responsible for the gender pronoun errors in L2 English ("pro-drop hypothesis"). She conducted an experiment eliciting semi-spontaneous speech in English, where she compared native Spanish and native French speakers of L2 English with respect to the pronoun errors they produced. French was chosen as it is a Romance language that is similar to Spanish in several aspects, but which, in contrast to the Spanish language, is not a pro-drop language. Each test group consisted of 20 participants who were comparable in terms of education, age of English acquisition, frequency of use and proficiency. The participants were shown 43 illustrations and were asked questions designed to elicit pronoun production. The subjects were instructed to respond freely, and the pronoun errors they produced were recorded. The types of reported errors fall in the following categories: person errors (e.g., 'I' instead of 'you'), number errors (e.g., 'I' instead of 'we'), gender errors ('he' instead of 'she' and vice-versa), animacy errors (e.g., 'he' instead of 'it'), omission errors (e.g., 'is swimming'), insertion errors (e.g., 'the boy he played' instead of 'the boy played') and other errors (e.g., 'it' instead of 'there' in 'there is').

Spanish speakers of L2 English indeed made significantly more gender errors (4.30%) compared to other types of pronoun errors and to the French group (0.68%)[1]. The pronoun errors recorded were not due to erroneous transfer of the Spanish L1 grammar, as the Spanish speakers made no omission errors ('is swimming'); thus, in none of the items of Antón-Méndez's experiment did the subjects omit a pronoun, which would have been the case in a grammatical transfer. Importantly, even though there were slightly more 'he' than 'she' errors (he: 5.68%, she: 2.98%), the difference is not statistically significant. Therefore, the Spanish speakers were not using a default pronoun (e.g., always 'he' instead of 'she'). The use of 'he' as the default pronoun would have suggested that another factor might underlie the error, for instance, the difficulty that the English phonology poses for speakers of Spanish. The Spanish phonology does not contain the phonemes /ʃ/ in 'she' and /h/ in 'he', therefore one explanation for the gender pronoun issue could be at the phonological level. In the present study we focused only on the pro-drop feature, not because we disregard the potential role of the phonology, but because we wanted to investigate whether the pro-drop feature has the capacity of causing this type of gender errors in L2.

In order to focus on the pro-drop feature, we simulated bilingual sentence production using computational cognitive modeling. The pro-drop feature is not the sole difference between the French and Spanish languages, and one could argue that the differences in the error patterns between the two groups could have been partially attributed to confounding factors, for instance, to a different L2 English teaching system in Spain and France.

Using computational modeling we can remove all possible confounds and therefore minimize the variance by focusing only on the phenomenon of interest, which in this case is the pro-drop feature and its possible effect on L2 English pronouns. For this reason, we modified Dual-path (Chang, 2002), a computational cognitive model of sentence production, to account for bilingualism. We then compared L2 English speech production of simulated native speakers of Spanish (ES) on the one hand, to L2 production of simulated native speakers of a Spanish-like language ('non-pro-drop Spanish', NPD_ES) on the other hand. The latter contained all the features of the Spanish language (lexicon, allowed structures) except the pro-drop feature; therefore, pronouns needed to be used at all times. All input languages (ES, NPD_ES and EN) were artificially generated and based on the Spanish and English language, using a subset of their lexica and syntactic structures. If the bilingual Spanish-English (ES-EN) Dual-path model produces significantly more subject pronoun errors in English than its Spanish-like non-pro-drop equivalent (NPD_ES-EN), it will be clear that the pro-drop feature of the Spanish language is the reason for this particular L2 error in the simulation, as the two simulated languages differ only in their pro-dropness. If this is the case, we will have confirmed

---

[1]The percentages are calculated by Antón-Méndez (p. 129, Table 6) and they represent the frequency of the gender pronoun mistake (68 and 10, respectively) with respect to the total number of pronouns produced where this particular mistake could have occurred.

that the pro-drop feature has the capacity to lead to gender pronoun errors in L2 English.

Method

In order to simulate Spanish speakers of L2 English, we developed two bilingual models using a modified version of Dual-path which is a connectionist model based on the Simple Recurrent Network (SRN; Elman, 1990) achitecture (Chang, 2002).

Bilingual Dual-path model

Dual-path (Figure 1) learns to convert a message into a sentence by predicting the sentence word by word ("next word prediction"). It has two pathways (hence the name) that influence the production of each word; the meaning system which learns concepts, roles and event semantics, and the sequencing system which is an SRN that learns to abstract syntactic patterns. Both paths influence the word output layer. The sequencing system consists of one recurrent hidden layer (of 30 units in our simulations) and two "compress" layers (of 12 units each) that are placed between the input word, the hidden layer and the output word.

The meaning system learns to map the input word onto a concept, which is linked to a specific thematic role (that is given for each sentence through fixed connections). The fixed connections allow the separation between concepts and roles, which, in turn, enables the model to generalize and to produce words in novel places. The thematic role is connected to the hidden layer, and so is the "event-semantics" layer. The hidden layer spreads the activation to the next thematic role (in the meaning path, and the "compress" unit in the syntactic path), which is in turn linked to a specific predicted concept that is used as input to the output word layer, along with the "compress" unit.

In the original model, all layers use the tanh activation function, except the output layer that uses softmax. In the modified version of the model, we also employed softmax for the predicted role layer. This led to a stricter selection of the upcoming thematic role which helped overcome a difficulty that the model had with learning the correct articles regarding gender and definiteness (e.g., 'a' vs 'the'). Furthermore, our version has a "target language" layer in the meaning path that is used as an additional input to the hidden layer, along with the "event-semantics" layer. The "target language" denotes the intended spoken language and helps the model handle more than one language. The modified model can be found at https://github.com/xtsoukala/dual_path .

Input languages

Message Dual-path is trained using randomly generated sentences paired with their meaning (Chang, Dell, & Bock, 2006). The meaning (message) contained information regarding four thematic roles (AGENT, PATIENT, ACTION, RECIPIENT). A concept (e.g., 'WOMAN' for the English word 'woman' or Spanish word 'mujer') was assigned to each thematic role depending on the meaning that needed

to be expressed (e.g., in the sentence "the woman run -s" the message would include AGENT=WOMAN, DEF). Furthermore, the message contained event-semantic information (denoted as 'E'), which gave information regarding the tense (PRESENT or PAST) and aspect (SIMPLE or PROGRESSIVE). The message contained information about the target language (ES or EN) as well. This information was given at the beginning of the sentence along with the roles and the event-semantics, so that the model knew whether it was supposed to produce an English or Spanish sentence.

Structures The allowed structures for all languages were the following (where 'S', the subject, is omitted in the pro-drop case):

1. (S)V: (Subject) - Verb, e.g., "He runs"

2. (S)VO: (Subject) - Verb - Object, e.g., "She kicked the ball"

3. (S)VIODO: (Subject) - Verb - Indirect Object - Direct Object, e.g., "He gave the girl a book"

4. (S)VDOIO: (Subject) - Verb - Indirect Object - Direct Object, e.g., "He gave a book to the girl"

The sentences in English and in non-pro-drop Spanish always started with a pronoun, and the sentences in pro-drop Spanish never started with a pronoun but always with a verb.

Lexicon The total lexicon consisted of 34 nouns: 11 male ('man', 'boy', 'father', 'brother', 'dog', 'hombre', 'niño', 'padre', 'hermano', 'perro', 'maestro'), 11 female ('woman', 'girl', 'mother', 'sister', 'cat', 'mujer', 'niña', 'madre', 'hermana', 'gata', 'enfermera') and 12 inanimate ('ball', 'stick', 'toy', 'kite', 'key', 'bag', 'pelota', 'palo', 'juguete', 'cometa', 'llave', 'bolso'), 24 verbs (e.g., 'give', 'show', 'walk', 'throw', 'present', 'dar', 'lanzar', 'presentar', 'nadar', 'caminar') and 26 function words (e.g., articles ('a', 'the', 'un', 'una', 'el', 'la'), pronouns ('he', 'she', 'él', 'ella') and auxiliary verbs ('is', 'was', 'está', 'estaba')).

The model treats the verb lemma ('give') and the suffix ('-s') as two different units. Note that syntactic information (such as 'verb', 'noun') is not given explicitly, but is learned by the model during training through the syntactic path. The syntactic gender was also learned implicitly during training through the article of Noun Phrases (NP) and pronouns. Semantic gender (e.g., 'ACTRESS, F', 'ACTOR, M') was not included in the model.

Thematic roles could be expressed using either an NP with definite (DEF) or indefinite (INDEF) articles (e.g., 'the woman', 'a woman') or the pronoun (PRON) equivalent ('she').
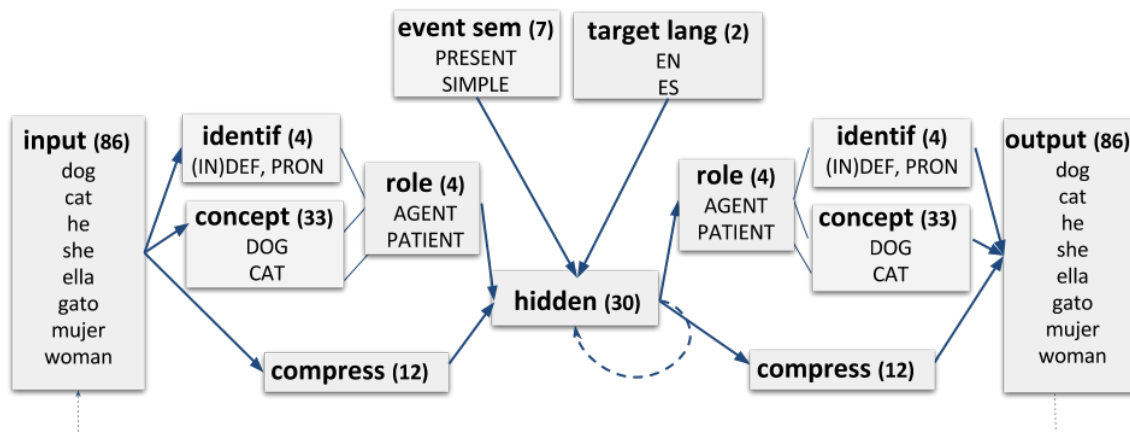
Figure 1: Bilingual Dual-path model

Example The following message would be the same across languages:

    AGENT=WOMAN, PRON;

    ACTION=GIVE;

    PATIENT=INDEF, KEY;

    RECIPIENT=DEF, GIRL;

    E=SIMPLE, PRESENT, AGENT, PATIENT, RECIPIENT

and it would be expressed linguistically in the following manner for the three languages:

1. she give -s the girl a key . [EN]

2. d -a a la niña una llave . [ES]

3. ella d -a a la niña una llave . [NPD_ES]

Training

The two models were trained on 2000 randomly generated sentences (training set) and tested on 500 unseen sentences (test set). The models contained almost identical sets, with the only difference that the NPD_ES model expressed the subject pronoun at all times, whereas the ES model never did and always started with a verb. For each model we ran 100 simulations using the same input, but different random initial weights per simulation, as the input and the weights are the only non-deterministic parts of the model. The models were trained for 20 epochs, where 1 epoch corresponds to a full iteration of the training set (2000 sentences). At the beginning of each epoch, the training set was shuffled. In order to simulate late L2 acquisition, we first trained the models for 20 epochs using Spanish input only, and then used the fully trained weights as initial weights for the bilingual models. The bilingual input consisted of newly generated (2000 training and 500 test) sentences, this time using 50% (pro-drop or non-pro-drop according to the model) Spanish and 50% English. We excluded from the analysis 7 simulations that did not manage to learn at least 75% of the test set by


Figure 2: Performance on the training and test sets over the training period (20 epochs) averaged over 93 simulations for the two bilingual models. Performance is measured in percentage of correctly produced Spanish and English sentences.

the end of the training in one of the two models, leading to a total of 93 simulations. Both bilingual models were able to perform equally well by the end of the training, reaching 99.69% correct for ES-EN and 99.70% correct for NPD_ES-EN (Figure 2) on the test set that contained English and Spanish sentences.

Results

In order to assess the performance of the two bilingual models on L2 pronouns, we focused only on the English sentences (50% of the test set). If a pronoun error was detected and the sentence was grammatical, it was classified as a gender pronoun error. We compared the performance of the two bilingual models with regard to the gender pronoun error production. If the models had a comparable performance we would not be able to confirm that the pro-drop feature has the capacity to lead to gender pronoun errors in L2 English. If, on the other hand, the NPD_ES model made fewer gender pronoun

Figure 3: Production of English gender errors (in log scale) averaged over 93 simulations for the two bilingual models. Note that the gaps in the non-pro-drop model are because log(0), for 0% error rate, is not defined, and therefore not plotted.

errors than the ES model it would indicate that the pro-drop feature is a possible explanation.

The non-pro-drop Spanish-English (NPD_ES-EN) bilingual model (Figure 3) produced almost no gender pronoun errors (maximum percentage: 0.11%) whereas the bilingual model based on pro-drop Spanish (ES-EN) initially produced 9.75% pronoun errors, gradually dropping to 0.05%.

Crucially, the ES-EN model never reached 0% (minimum error rate: 0.02%) whereas the NPD_ES-EN model did. Following visual inspection, we ran a z-test for proportions from epoch 5 onwards to test for a difference in error rate between the models. The difference is significant (z=7; p<.001).

Discussion

Our simulations showed that a bilingual model with L1 pro-drop Spanish and L2 English produced significantly more gender pronoun errors than a similar model with L1 non-pro-drop Spanish. These sentences were grammatically correct: the only error they contained was a pronoun with incorrect gender. Given that the only difference between the two L1s was the pro-drop feature, we have demonstrated that the pro-drop nature of Spanish can indeed cause the gender pronoun error as observed in L1 Spanish speakers of L2 English.

Why the pro-drop feature does not lead to a direct language transfer ("is walking") in either the model or humans remains to be investigated, as the current simulations and results do not explain how pro-dropness in L1 could lead to gender errors in L2. Nevertheless, having a computational model that simulates the gender pronoun errors in L2 English can point us in the right direction. Our hypothesis for the occurrence of the gender error is that the gender information is not as crucial for the message planning, at least in the subject position, of a pro-drop language, and is therefore weaker or omitted,

even when producing sentences in a non-pro-drop L2.

It is important to point out that the Dual-path model does not contain a phonological level (Garrett, 1988). One might have thought that the reason Spanish speakers confuse the words 'he' and 'she' is because of the difficulty the English phonology poses for native speakers of Spanish. However, our simulations have produced gender errors without having any phonological representations. This does not mean that phonology could not play a role, but rather that it is not the only possible explanation.

It is also crucial to note two simplifying assumptions in these simulations. First, as mentioned in the Method section, the input for all three languages (EN, ES, NPD_ES) was artificially generated and it only represented a subset of the actual languages. In general, using natural input would be preferable as it would increase the validity and naturalness of the results. However, the benefit of miniature languages that are typically used in cognitive modeling is that they can be easily manipulated. For instance, in the simulations described here we were able to add and remove the pro-drop feature at will, leaving everything else the same, and thus to isolate this important feature from confounding factors.

Second, a crucial simplifying assumption in the miniature language is the absence of full NP subjects. We therefore repeated the simulations using new input for all languages, this time including 50% pronouns at the subject position and 50% noun phrases. Preliminary simulations show no gender errors in either model, which means that further research is needed using more natural language input, starting with a more naturalistic proportion of pronouns and NPs in the subject position based on English and Spanish corpora.

Conclusion

Computational modeling can be used to validate or generate linguistic hypotheses while focusing on specific factors of interest and minimizing the variance. In this study, we have addressed the question as to whether the pro-drop feature of the Spanish language has the capacity to cause the gender pronoun errors that Spanish speakers of L2 English have been shown to produce (Lahoz, 1991; Antón-Méndez, 2010). The reported simulations showed that the model with L1 pro-drop Spanish produced more gender pronoun errors in L2 English than the model with L1 non-pro-drop Spanish, which is a necessary but not sufficient condition for the pro-drop hypothesis.

Acknowledgements

References

Antón-Méndez, I. (2010). Gender bender: Gender errors in L2 pronoun production. *Journal of Psycholinguistic Research*, *39*, 119–139.

Anton-Mendez, I. (2011). Whose? L2-English speakers' possessive pronoun gender errors. *Bilingualism: Language and Cognition*, *14*, 318–331.

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

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

Davidson, B. (1996). 'Pragmatic weight' and Spanish subject pronouns: The pragmatic and discourse uses of 'tú' and 'yo' in spoken Madrid Spanish. *Journal of Pragmatics*, *26*, 543–565.

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

Garrett, M. F. (1988). Processes in language production. *Linguistics: The Cambridge survey*, *3*, 69–96.

Lahoz, C. M. (1991). Why are he and she a problem for Spanish learners of English? *Revista Española de Lingüística Aplicada*, 129–136.

Odlin, T. (1989). *Language Transfer: Cross-Linguistic Influence in Language Learning*. Cambridge: CUP.

Poulisse, N. (1999). *Slips of the tongue: Speech errors in first and second language production* (Vol. 20). Amsterdam: John Benjamins.

White, J., Muñoz, C., & Collins, L. (2007). The his/her challenge: Making progress in a 'regular' L2 programme. *Language Awareness*, *16*, 278–299.