Emma Chabane

A noisy channel framework for sentence comprehension in Spanish and French

Linguistic communication is no easy task, and is very often subject to noise : speakers can make errors or introduce ambiguity, and our environment can introduce other sounds which might make it difficult to comprehend one another. Nonetheless, people can generally understand each other easily, therefore suggesting that the way humans process and understand language accounts for the omnipresent noise in our lives.
As such, one can make inferences about the speaker's intended meaning when utterances have been subjected to potential noise. Given prior information and a noise model, that is, what they know about the world, people can make inferences to understand what they perceive.
According to this model, the noisy channel framework, the probability $P(S_i|S_p)$ of the perceived sentence(S_p) being interpreted as the intended sentence (S_i) is proportional to the prior probability of the intended sentence given a person's world knowledge as well as the probability that the intended sentence could have been mistakenly transformed to the perceived sentence during linguistic communication:

$$P(S_i|S_p) \propto P(S_p|S_i) * P(S_i)$$

Since the introduction of the noisy channel framework (Levy, 2008), various studies have aimed to test the model with different stimuli and variants. Gibson et al (2013) showed that perceived linguistic utterances act as a starting point for sentence comprehension: given syntactically correct but semantically implausible sentences such as "the mother gave the candle the daughter," participants infer a plausible meaning about half the time, and interpret the perceived sentence as one that has been corrupted. People indeed use their world and previous linguistic knowledge to give meaning to utterances that don't immediately make sense. The literature on the noisy channel framework has helped investigate the nature of the framework using English data, and refine researchers' understanding of how the noise model functions in the English language . Gibson, Bergen & Piantadosi's 2013 paper first suggested that the noise model should include deletions and insertions of function words from an intended sentence, and hypothesized that a corrupted sentence caused by deletions is much more likely than one caused by insertions. Poppels & Levy (2016) built on top of these latter results and published crucial research which put forward the idea of exchanges of category matched words as a part of the noise model. More recently, Ryskin et al (2018) reviewed those three potential sources of noise and proposed that deletions and exchanges from the intended sentence were more likely than insertions of a word.

As such, the noise model has been studied extensively in English. However, the noise model itself has only been tested on a small set of languages other than English. A few other papers (in languages other than English) do not focus on the noisy model part of the noisy channel framework and instead study one's world knowledge and its impact on interpretation. In other words, how prior sentence frequency and probability affects interpretation. Liu et al's (2020) research on Mandarin Chinese for instance, confirms the robustness of one main component of the noisy channel framework: prior information. Mandarin speakers indeed interpret sentences based on the probability of certain sentence structures. More recently, papers have started looking into and testing the noisy model itself in other languages. This was done in Zhan et al (in review) on Mandarin speakers, which revealed that participants inferred a plausible

meaning from an implausible sentence the most when the intended sentence had been corrupted by a deletion or insertion; inferences were not as likely when sentences were corrupted by exchanges. This is in line with past literature on English data. Another paper, Poliak et al (in press) focused on the meaning prior and looked into how Russian speakers interpret sentences when canonicality changes. This paper showed Russian speakers use the structural prior in order to interpret the implausible utterances given to them.

Fleming et al's 2022 paper endeavors to test the framework on a new language, Spanish. This 9.59 final project aims to replicate the study done by Fleming et al, which looked into the robustness of the noisy channel model when studying sentence interpretation in Spanish. Following this replication, I will extend the research to testing the noisy channel model on the French language.

Methods

Experiment 1 : Replication of Fleming et al's noisy channel paper for Spanish

This experiment closely followed the methodology and design used in Fleming et al. (2022). The subjects recruited for this study were a total of 60 native Spanish speakers from all around the world (Average age of exposure = 0.58 years, Average years of exposure = 27.01 years) on the crowdsourcing website Prolific. Contrary to Fleming et al, which ran one study on Mexican Spanish speakers and a subsequent one on Castilian Spanish speakers, we chose to give our study to Spanish speakers regardless of the dialect. Indeed, dialectical differences did not provide an explanation for the results from Fleming's study. Participants were presented with syntactically correct sentences that were either plausible or implausible. Following a 2 by 2 design, we ran a first experiment with active or passive sentences, and a second with transitive and intransitive sentences, as in *Table 1*.

	Active	**Passive**	**Transitive**	**Intransitive**
Plausible	El chef ha cocinado el bistec (The chef has cooked the steak)	El bistec fue cocinado por el chef (The steak was cooked by the chef)	La olla cocía la comida (The pot cooks the food)	La comida cocía en la olla (The food cooks in the pot)
Implausible	El bistec ha cocinado el chef (The steak has cooked the chef)	El chef fue cocinado por el bistec (The chef was cooked by the steak)	La comida cocía la olla (The food cooks the pot)	La olla cocía en la comida (The pot cooks in the food)

Table 1: Example stimuli provided to participants for a replication of Fleming et al's study. Sentences were written by the team working with Fleming et al.

Both the stimuli and fillers in the replication were extracted from Fleming's own study. Implausible sentences in both the active/passive and transitive/intransitive forms were corrupted based on the three main components of the noisy channel model that has been researched up to this point: active and passive sentences were corrupted by exchanges, transitive sentences were

corrupted by deletions from an intransitive sentence, and intransitive sentences were corrupted by insertions from a transitive sentence.
Participants answered a yes or no comprehension question following the sentence that they just read. The way participants answered the comprehension questions gauged whether they interpreted said sentences literally or not, that is, whether or not they made inferences about the intended meaning of the sentences, given the perceived sentence.

Experiment 2: Extension of the noisy model to French

The second experiment aimed to test the noisy channel framework in French and extend Fleming's work to a language similar to Spanish to explore cross-linguistic differences of the noisy channel framework. Native French speakers (Average years of exposure = 32.77) were recruited on Prolific for this study. Participants were by and large French and Canadian French language speakers. Participants underwent the same procedure as the native Spanish speakers. That is, they read both implausible and plausible sentences in two possible constrictions: active/passive and transitive/intransitive. Subsequently, participants answered a yes or no comprehension question about the sentence they just read. The stimuli was developed and written by the author, and verified by a colleague, both of whom are native French speakers; Table 2 presents an overview of example stimuli provided in the experiment. As such, the sentences were tailored to the grammatical structure of French, and written such that all plausible sentences seemed intuitively acceptable to a native French speaker.

	Active	**Passive**	**Transitive**	**Intransitive**
Plausible	L'espion a infiltré l'agence. *The spy infiltrated the agency*	L'espion a infiltré l'agence. *The spy infiltrated the agency*	L'eau de Javel blanchit le t-shirt. *The bleach whitens the shirt.*	L'eau de Javel blanchit le t-shirt. *The bleach whitens the shirt.*
Implausible	L'agence a infiltré l'espion. *The agency infiltrated the spy.*	L'agence a infiltré l'espion. *The agency infiltrated the spy.*	Le t-shirt blanchit l'eau de Javel. *The shirt whitens the bleach.*	Le t-shirt blanchit l'eau de Javel. *The shirt whitens the bleach.*

Table 1: Example stimuli provided to participants in experiment 2. Sentences were written by the author working with on the extension to French.

Creating the transitive/intransitive sentences in French was the bulk of the work, and required some careful thought about how to develop plausible sentences that make sense and that can be corrupted such that participants could infer the intended sentences. Compared to English and Spanish, it was much harder to find constructions that would work for all 4 sentences that needed to be created in the 2 by 2 transitive/intransitive and plausible/implausible design. There are 2 reasons for that. One, depending on the subject of a verb, in French the verb might switch to its reflexive form when going from transitive to intransitive. Two, switching from an intransitive form to a transitive form of the sentence could transform the verb construction by using a

causative alternance, that is, add the causative “faire” (“do”) before the main action verb. These two potential issues made it complicated to write stimuli that worked well. Nonetheless, the author spent a lot of time writing the sentences, and used ergative verbs, which can be used in both their transitive and intransitive forms. Although they are less frequent in French compared to languages like English, certain verbs allowed for the straightforward creation of the stimuli; this includes verbs describing color, state, and temperature changes. For the extension, twenty items for each active/passive and transitive/intransitive constructions were created. Half of the fillers were adapted from the Spanish replication, while the other half of the fillers came from a completely unrelated study from a colleague.

Predictions

Experiment 1 : Replication of Fleming et al’s noisy channel paper for Spanish

We predict that we will achieve similar results as Fleming et al in our replication of their study. Indeed, Fleming’s results were significant enough that it would seem unlikely for us to get results that would go in a totally different direction. As such, we predict that, one, active and passive sentences will generally be interpreted literally more often than transitive and intransitive sentences, meaning exchanges are seen as less likely than deletions and insertions. Two, we expect to see similar trends in the interpretation of implausible intransitive sentences given past results: specifically, we expect sentences following that construction to elicit much more inferences than any other.

Experiment 2: Extension of the noisy model to French

Given the linguistic similarities between French and Spanish, we predict that our results will fall in line with what was found in Spanish. That is, we expect to see intransitive implausible sentences eliciting more inferences compared to other constructions, which means that participants can infer the intended meaning of the sentences, even if these sentences are corrupted by insertions. We also expect active and passive sentences to be interpreted literally more often than the intransitive/transitive form.

Results

For analysis, the authors used RStudio as well as the TidyR and ggplot libraries on the table data and results extracted from MIT Surveyor.

Experiment 1 : Replication of Fleming et al’s noisy channel paper for Spanish

Our replication of Fleming et al’s paper yielded similar results to what had initially been found in the study. Most notably, intransitive implausible sentences, once again and by far, resulted in much more inferences than any other construction. 69.6% of intransitive implausible sentences were interpreted non-literally, that is, participants inferred a plausible meaning more than half of the time for this construction.

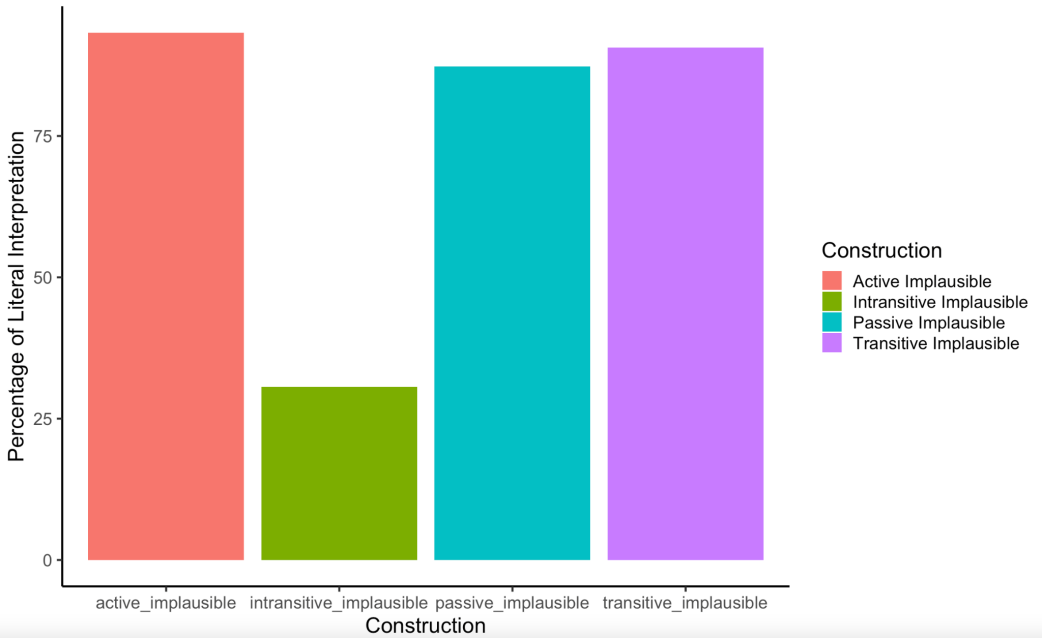

Figure 1: Bar plot presenting the percentage of literal interpretation in each Spanish sentence construction given to native speakers of the language. Intransitive implausible sentences were interpreted literally around 30% of the time.

Moreover, transitive sentences are interpreted literally slightly more than passive sentences, although this might not be a significant result. Overall, the results show that active, transitive, and passive sentences have a high, near-ceiling rate of literal interpretation, while intransitive implausible sentences elicit a very high number of inferences.

Experiment 2: Extension of the noisy model to French

Our test of the noisy channel model in French provided us with results that were in line with our expectations as well as Fleming's previous results. Although not as significant as our replication, the rate of literal interpretation of intransitive implausible sentences was much lower still than any other constructions; intransitive implausible, caused by insertions, elicit inferences about 52% of the time. Active implausible sentences, caused by exchanges, were interpreted more than 90% of the time. Lastly, passive implausible and transitive implausible had similar rates of literal interpretation, around 75 to 80%. Plausible sentences, although only provided for comparison, all had near-ceiling interpretation rates. Therefore, there is a significantly lower literal interpretation for intransitive sentences, in line with our expectations.

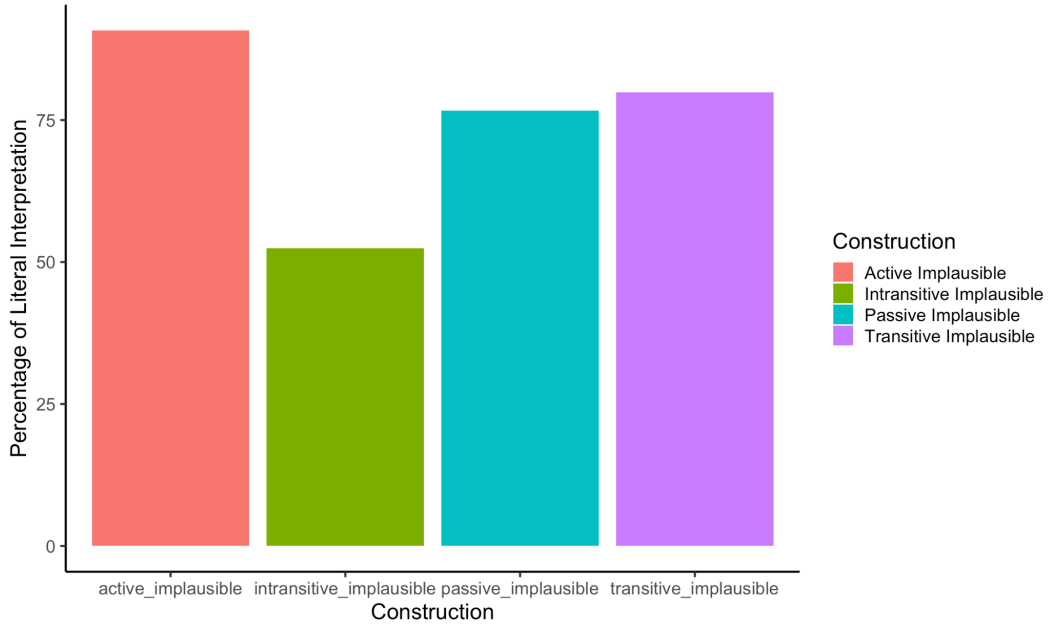

Figure 2: Bar plot presenting the percentage of literal interpretation in each French sentence construction given to native speakers of the language. Intransitive implausible sentences were interpreted literally around 58% of the time, while other constrictions had a 75% or above of literal interpretation.

Discussion

Both the Spanish replication and the French extension yielded our expected results, given what was previously found in Fleming et al. Specifically, for both languages, our key finding is that intransitive implausible sentences are interpreted non-literally much more often than any other constructions used in this study. This result, which has been confirmed in both of these studies as well as Fleming's research, shows that transitive and intransitive sentences may be interpreted differently in French and Spanish, than in English.

This makes sense, given that corpus frequencies, and verb frame frequencies may easily change from language to language. Cross-linguistic differences make it such that differences in the verb frame frequencies, syntactic structures or language-specific alterations may also impact the patterns of sentence interpretation in these languages.

Our replication of Flemig's study aimed to assess the validity of the results found by the authors, and evaluate the noisy channel framework in Spanish. This replication examined if the observed results and patterns of sentence interpretation are consistent for different groups of Spanish speakers, especially given Fleming's results on the interpretation of intransitive implausible sentences. Our results suggest that this result is robust, and therefore provided more evidence of the generalizability of the previous findings. These results give valuable insights into the application of the noisy channel framework to other languages and its robustness across cross-linguistic differences.

Extending this research to the French language has allowed us to explore more about cross-linguistic differences in relation to the noisy channel model. French was not only chosen as

it is the author's first language, but also because it is a language that is similar to Spanish in many ways. Because of this, it was expected that certain similarities would emerge in terms of sentence comprehension patterns for both languages. The extension to French confirms the similarity in sentence interpretation between French and Spanish, and provides crucial information in our understanding of how linguistic differences may affect the human language interpretation process.

However, we cannot say yet what these results say about the cross-linguistic robustness of the noisy channel model, even if these results do not fall in line with what was previously found on English data in the framework. Specifically, previous results had shown that intransitive implausible sentences were interpreted literally more often than their transitive counterpart (Gibson et al, 2013). Nonetheless, other cross-linguistic factors may affect sentence interpretation in French and Spanish. Although this research has led to interesting results, simply asking yes/no comprehension questions is a bit of a black box. It does not truly explain what goes on in participants' processing of the language, and further research should look into sentence interpretation in these languages with a design that might be more telling of what the participants infer or do not infer.

Interestingly, in both languages as well, transitive sentences were interpreted literally as often as active and passive sentences, even surpassing passive sentences. This is an interesting result as transitive implausible are, in English, explained by a deletion, which is an edit type more likely than insertions and exchanges, per past literature. In the Spanish experiment (and in other studies in English), Fleming predicted implausible transitive sentences would elicit more inferences given that they are corrupted by deletions (compared to insertions for the intransitive). Nevertheless, a potential alternative explanation, which has been mentioned in past literature (Gibson, Bergen, & Piantadosi, 2013), is that the transitive material could also involve an exchange:

A. Le t-shirt blanchit l'eau de Javel. (transitive implausible)

a. "The shirt whitens the bleach."

B. Le t-shirt blanchit dans l'eau de Javel. (intransitive plausible)

a. "The shirt whitens in the bleach."

C. L'eau de Javel blanchit le t-shirt. (transitive plausible)

a. "The bleach whitens the shirt."

Sentence B could indeed be corrupted by a deletion and result in the implausible sentence A, or alternatively, sentence C could be corrupted by an exchange and result in sentence A. If sentence C is more likely given the French language prior (ie. the transitive version is more often used), it is possible that participants make inferences less often in the transitive implausible case, given that the most likely intended meaning was caused by an exchange, which has been shown to be less likely than insertions or deletions.

In the future, research should look into replicating my French study with a bigger sample size of participants to improve the generalizability of the findings. Lastly, a main next step in this research would be to do a corpus analysis of French and Spanish to figure out how the linguistic

prior in these languages may affect sentence interpretation, and whether certain sentence and verb constructions are more likely. This analysis would provide us with crucial information to better interpret and understand the results we've received.

Conclusion

The present replication of Fleming et al's study as well as our own extension to the French language has provided us with valuable insights into sentence interpretations of these two languages. The replication provided confirmation of the previously found results, and our extension gave us excellent support for our previous hypothesis, as well as novel insights on sentence comprehension in the noisy French. These studies provide interesting information about the applicability of the noisy channel framework to other languages and shed light on cross-linguistic differences of the interpretation of utterances under noisy input. By examining prior language-specific information through a corpus analysis, such as the verb frame frequencies of French and Spanish, this research has the potential to enhance our understanding of the human language comprehension process, and to ultimately help develop technology that understands language like we do. Further research is nonetheless necessary to fully grasp the nuances of cross-linguistic interpretation within a noisy channel framework.