

SPACER: A Parallel Dataset of Speech Production And Comprehension of Error Repairs

Shiva Upadhye* Jiaxuan Li* Richard Futrell
{shiva.upadhye, jiaxuan.li, rfutrell}@uci.edu
Department of Language Science
University of California, Irvine
**equal contribution*

Abstract

Speech errors are a natural part of communication, yet they rarely lead to complete communicative failure because both speakers and comprehenders can detect and correct errors. Although prior research has examined error monitoring and correction in production and comprehension separately, integrated investigation of both systems has been impeded by the scarcity of parallel data. In this study, we present *SPACER*, a parallel dataset that captures how naturalistic speech errors are corrected by both speakers and comprehenders. We focus on single-word substitution errors extracted from the Switchboard corpus, accompanied by speaker’s self-repairs and comprehenders’ responses from an offline text-editing experiment. Our exploratory analysis suggests asymmetries in error correction strategies: speakers are more likely to repair errors that introduce greater semantic and phonemic deviations, whereas comprehenders tend to correct errors that are phonemically similar to more plausible alternatives or do not fit into prior contexts. Our dataset¹ enables future research on integrated approaches toward studying language production and comprehension.

1 Introduction

Production errors are common in naturalistic speech; however, they rarely lead to a complete breakdown in communication, as interlocutors are able to monitor, detect, and repair errors in real-time. For this reason, characterizing the process of error correction has remained a shared goal of both language comprehension and production research.

Comprehenders process errors by integrating perceived linguistic input with prior context and expectations, and might arrive at an interpretation different from the literal meaning of the linguistic input (Ferreira et al., 2002; Ferreira and Stacey,

2000; Dempsey et al., 2023; Bader and Meng, 2018; Levy et al., 2009; Gibson et al., 2013). One possible mechanism is that they perform rational inference over the perceived errors (Levy et al., 2009; Levy, 2008; Gibson et al., 2013; Futrell et al., 2020; Ryskin et al., 2018; Poppels and Levy, 2016; Zhang et al., 2023). When the perceived form is incongruent with prior context or similar to a more plausible alternative, comprehenders might override the literal input and reconstruct an alternative interpretation.

In language production research, speech errors have played a crucial role in shaping our understanding of the cognitive machinery of production, including the role of online control (Fromkin and Fromkin, 1973; Levelt, 1983; Dell, 1986). Numerous studies analyzing the temporal and distributional properties of speech errors have found evidence of a two-stage monitoring and correction process, which operates first on internal representations and then on the articulated linguistic signal (Levelt, 1983; Blackmer and Mitton, 1991; Hartsuiker and Kolk, 2001; Nooteboom and Quené, 2017). Although the mechanism of monitoring has remained a point of contention in the literature (Levelt, 1999; Nozari et al., 2011; Hickok et al., 2011; Roelofs, 2020; Gauvin and Hartsuiker, 2020), accounts of repair processing have posited sustained competition between activated representations and selection control as potential mechanisms of correction (Hartsuiker and Kolk, 2001; Nozari et al., 2016; Nooteboom and Quené, 2019; Gauvin and Hartsuiker, 2020).

Much of our understanding of speakers and comprehenders’ error correction strategies comes from research traditions that have made limited contact with each other. In particular, existing datasets or experimental paradigms focus solely on corrective behavior in the *absence* of an interlocutor. However, in a communicative context, speaker choices may exhibit a balance between ease of production

¹The dataset and code are available at: <https://github.com/goldengua/SPACER-CMCL>

and communicative efficiency (Ferreira and Dell, 2000; Jaeger and Levy, 2006; Jaeger, 2010; Koranda et al., 2018; Goldberg and Ferreira, 2022; Futrell, 2023). For example, speakers may preemptively hyperarticulate words to improve comprehensibility (Aylett and Turk, 2004; Arnold et al., 2012; Meinhardt et al., 2020) or modulate the acoustic characteristics of their speech in response to listener feedback (Pate and Goldwater, 2015; Buz et al., 2016). Comprehension can be talker-specific as well: comprehenders can tailor their expectations to the speaker (Ryskin et al., 2020) and adapt their error correction strategies accordingly (Futrell and Gibson, 2017; Brehm et al., 2019).

In this study, we present a parallel dataset of Speech Production and Comprehension Error Repairs (*SPACER*) that captures how naturalistic speech errors are corrected by speakers and comprehenders. First, we compiled a corpus of naturalistic utterances with single-word substitution errors and repairs, as well as utterances that are not corrected by speakers. These utterances were then presented to comprehenders in a web-based text editing experiment, where each case is annotated by four to six comprehenders. Our dataset contains 1056 instances of naturalistic speaker’s utterances as well as 5808 comprehenders’ responses to speaker’s initial utterances. We also provide an exploratory analysis on how well comprehender’s error correction behavior can be predicted by lexical properties of speaker’s error and repairs, and vice versa. Our results suggest that asymmetries between error correction behaviors by comprehenders and speakers might be related to interaction between the two modes. The dataset offers resources to build experimental and computational work that bridges comprehension and production and informs how interaction affects strategic cue weighting in error monitoring.

2 Dataset

We focus on how comprehenders engage with utterances that may or may not have been corrected by a speaker in their original context. Figure 1 illustrates the design of the dataset. We assume that both the speaker and the comprehender can perform error correction, and we remain agnostic to the mechanism that accomplishes this process. Suppose a speaker produces an initial utterance (... *people don’t want the heat of the northeast*), which may contain an error, or non-optimal choice of words.

The speaker can monitor the initial utterance and *may* correct *heat* into *cold*. After receiving the initial utterance, comprehenders engage in an interpretation process, where the final response might not be the same as the literal meaning of the initial utterance.

The development of this dataset involved a two-step process. First, we identified and extracted utterances with and without word substitutions and overt repairs from the Switchboard corpus (Godfrey et al., 1992) of spontaneous speech. Subsequently, these utterances were presented as stimuli to participants in a web-based correction experiment. Hence, for each utterance, this dataset provides a speaker’s final production along with a group of comprehenders’ annotations. The two stages of this process are detailed below.

2.1 Naturalistic Speaker Correction Data

We identify and extract stimuli for the correction experiment from Switchboard NXT annotations (Calhoun et al., 2010), a subset of the Switchboard corpus that provides gold-standard disfluency annotations generated by human raters. While words in the corpus are annotated as *fluent*, *reparandum*, and *repair*, the *reparandum* label encompasses a variety of disfluencies such as filled pauses, false starts, repetitions, and substitutions. We programmatically identify substitutions using the following criteria. First, we only consider utterances with an equal number of *reparandum* and *repair* annotations to filter out instances where the speaker’s utterance plan may have undergone structural revisions. Next, we focused on utterances where (i) the word labeled *reparandum* (bolded) was immediately followed by a **non-identical** word that was labeled *repair* (underlined; see 1) or where (ii) the speaker repeated the *reparandum* sequence almost verbatim except for a single-word change (2). Finally, we eliminated instances where the *reparandum* was either a filled pause, false start, repetition, or a contracted form².

1. I I think that might be **talking** referring to uh something kind of uh alternative to the draft you know
2. So until I see the entire quote old guard of the Soviet **military** of the Soviet government

²We eliminate instances of contractions such as *I’ve* since they are reduced forms of multiword expressions such as *I have*

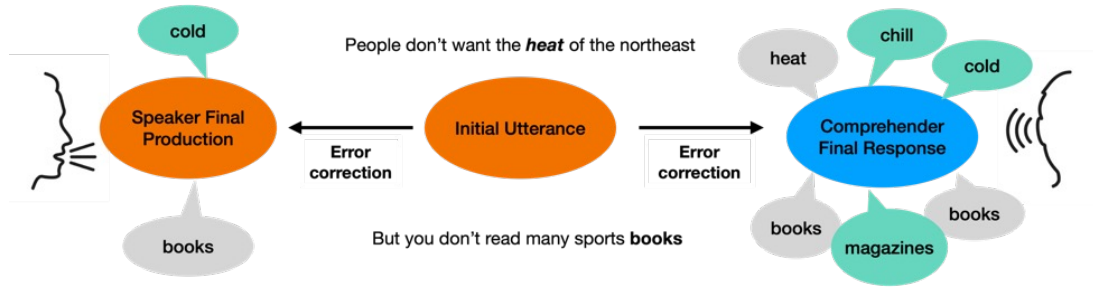


Figure 1: An illustration of the dataset design. Suppose a speaker produces an initial utterance. Both the speaker and the comprehender may engage in error monitoring and correction processes, resulting in speaker final production and comprehender final response being either the same as the initial utterance (represented in gray bubbles) or different from it (represented in green bubbles). Each utterance is annotated by four to six comprehenders.

completely roll over and disappear preferably buried I still consider them a threat

(see Appendix B for the detailed instructions provided to participants).

For utterances with multiple substitutions and repairs, we generate variants of the original utterance with each containing only one reparandum-repair position (see Appendix A for examples).

Finally, the selected utterances were screened for inappropriate content by three human annotators. This process yielded a total of 576 distinct utterance frames that featured errors repaired by the speaker (henceforth, *speaker corrected* or **SC** utterances). Furthermore, we also included 480 sentences from the same corpus that did not feature any overt repair made by the speaker (henceforth, *speaker uncorrected* or **SU** utterances; see Fig. 2).

2.2 Correction Experiment

We conducted human error correction experiments to understand how selected utterances are corrected during comprehension. 66 native English speakers participated in the experiment. Participants were recruited online via Prolific and compensated at \$16/hr. The experiment takes around 30 minutes.

The stimuli consisted of 1056 initial utterances, which included both the **SC** (*speaker corrected*) and **SU** (*speaker uncorrected*) utterances. We distributed the 1056 selected utterances into 12 lists. Each list contains 48 **SC** utterances and 40 **SU** utterances. Each list is annotated by four to six subjects.

The subjects are presented with initial utterances along with preceding context, and are instructed to check the quality of the last sentence from speech transcriptions and make necessary corrections by replacing the erroneous word with a more appropriate choice (Fig. 3). They also received explicit instruction to avoid insertion or deletion of word(s)

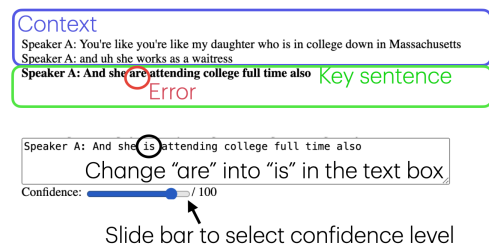


Figure 3: An illustration of comprehension experiment. A comprehender is presented with the *key sentence* together with preceding *context*. The comprehender is instructed to make necessary edits in the textbox and slide bar to indicate their confidence level.

We exclude subjects that made less than two corrections throughout the experiment ($N = 3$), and subjects who did not move confidence bars ($N = 3$). After subject exclusion, 528 trials were further removed because they contain word insertion or deletion, resulting in 5808 responses with either one or no substitution.

3 Analysis

3.1 Descriptive summary

As shown in Table 1, our dataset contain a total of 1056 initial utterances, where 576 initial utterances have been corrected by speakers (**SC** utterances). Each initial utterance is annotated by four to six comprehenders, yielding a total of 5808 responses with either one or no substitution. 34.7% of trials were corrected by the comprehender.

We focus on the items that are corrected by speakers (**SC**), and analyzed lexical properties of the critical target that has undergone correction (*error*) and its corresponding corrected form (*re-*

Speaker Uncorrected (SU)

Well when they meant the topic you know of what's wrong with the public schools my first impulse was it isn't really what's wrong with the public schools

Speaker Corrected (SC)

I don't think we've missed a fish store on the entire east → northeast coast of the United States

Like all these people are just standing → sitting around just going you know like what shall we do

Comprehender Uncorrected (CU)

Like all these people are just standing around just going you know like what shall we do

Comprehender Corrected (CC)

I don't think we've missed a fish → tourist store on the entire east coast of the United States

Well when they meant the topic you know of what's wrong with the public schools my first impulse → impression was it isn't really what's wrong with the public schools

Figure 2: Examples of *speaker uncorrected* (SU), *speaker corrected* (SC), *comprehender uncorrected* (CU), and *comprehender corrected* (CC) utterances in the SPACER dataset. Words highlighted in red were initially produced a speaker and later corrected. Words highlighted in green are corrections made by either the speaker in the original context or by a participant in the comprehension experiment. A grey highlight indicates that the word was not corrected by the participant. Note that for each SU and SC utterance, there may be up to four responses, which we classify as either CC or CU responses depending on whether or not the participant made a correction in their response.

	Comprehender		Total
	Corrected	Uncorrected	
Speaker Corrected	1437	1731	3168
Speaker Uncorrected	578	2062	2640
Total	2015	3793	5808

Table 1: Instances of *speaker corrected* and *speaker uncorrected* utterances that were corrected or remained uncorrected by participants in the correction experiment

pair). We first analyzed the part-of-speech (POS) categories of the critical words (error and repair). Figure 4 shows the distribution of part-of-speech categories for speaker initial produced errors and the corresponding repairs. The errors and repairs vary by POS category, with determiners (DET), pronouns (PRON), and verbs (VERB) exhibiting the highest frequency of both errors and repairs, whereas auxiliaries (AUX), particles (PART), and proper nouns (PROPN) show relatively fewer occurrences.

We further analyzed how the POS categories of speaker-produced errors would affect comprehender's error correction behavior. Figure 5 shows the number of corrected and uncorrected responses given speaker-produced errors with different part-of-speech (POS) categories. The results suggest variation in comprehenders' tendency to correct errors depending on the POS category of the presented errors, with higher correction rates for de-

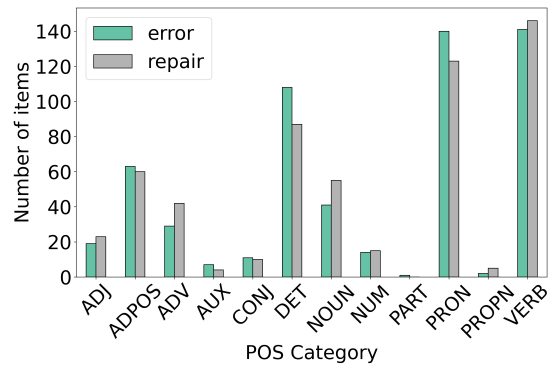


Figure 4: The POS categories of speaker-produced errors and corresponding repairs in *speaker corrected* utterances.

terminers (DET), verbs (VERB), and pronouns (PRON).

We examined how often comprehenders corrected errors in speaker-corrected versus speaker-uncorrected utterances. Figure 6 presents the proportion of corrected and uncorrected responses across items in the *speaker uncorrected* items (top panel) and *speaker corrected* items (bottom panel). A higher proportion of items were corrected in the *speaker corrected* items than in *speaker uncorrected* items. While initial utterances that are corrected by speakers are also more likely to be corrected by comprehenders, there is great variation between items.

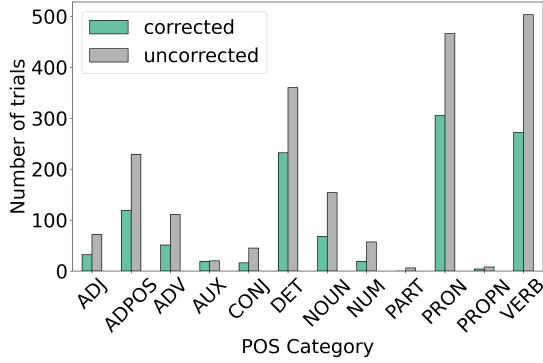


Figure 5: The number of corrected and uncorrected responses across different part-of-speech (POS) categories of presented errors in the *speaker corrected* utterances.




Figure 6: The proportion of corrected and uncorrected responses for each item in the *speaker corrected* condition (top) and *speaker uncorrected* condition (bottom). Each bar represents an individual item, with corrected responses shown in teal and uncorrected in gray.

3.2 Inter-rater agreement

Figure 7 shows the number of corrected and uncorrected responses by each subject. There is variability across subjects, with some participants correcting a substantial proportion of responses while others made few or no corrections. The average Cohen’s Kappa across all lists is 0.213, indicating fair agreement between subjects. Figure 8 displays pairwise agreement scores between subjects across all items (distributed in 12 lists). This suggests

that while some systematic agreement is present, variability in responses remains substantial.

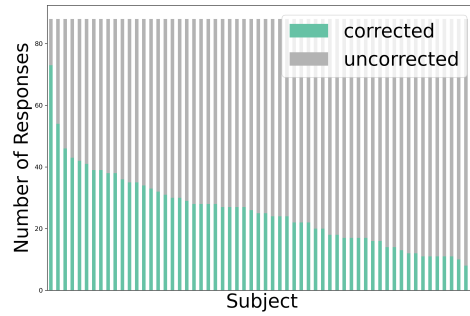


Figure 7: The number of corrected and uncorrected responses for each subject, sorted by the number of corrected responses.

3.3 Computational Metrics

For words that are corrected by either a speaker or a comprehender and their corresponding corrected counterparts, we calculated the following computational metrics: word predictability, word frequency in SUBTLEXus (Brysbaert and New, 2009), word length, semantic and phonemic distance between the initial word and the corrected word. We define word predictability as the log probability of the target word given preceding context, which we estimated from a pre-trained GPT-2 (small, uncased) transformer language model (Radford et al., 2019; Misra, 2022). The semantic distance is calculated using the cosine distance between GloVe word embeddings (Pennington et al., 2014). To compute the phonemic distance, we first convert the words to their corresponding IPA forms using the Carnegie Mellon Pronouncing Dictionary, and subsequently compute the Levenshtein edit distance between these forms.

4 Exploratory Analysis: Modeling Corrections

The dataset could provide opportunities for the investigation of asymmetric strategies on error correction during language production and comprehension. We present an exploratory analysis that examines the parallels between speaker and comprehender corrections. In particular, we ask (i) whether the lexical properties of the speaker’s error and repair are predictive of a comprehender’s decision to correct and (ii) whether comprehender’s responses over the perceived initial utterance are

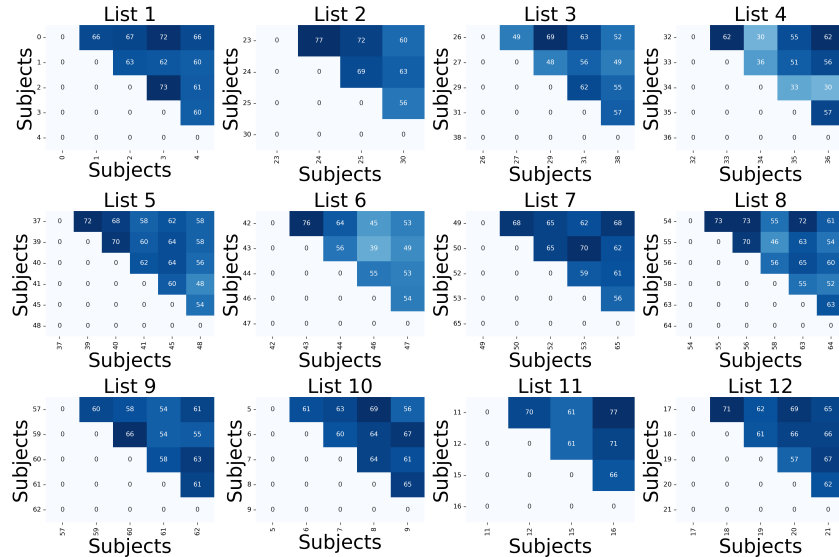


Figure 8: Pairwise agreement scores between subjects across 12 different lists of items.

predictive of the speaker’s decision to repair. We model both these choices as functions of the relative **frequency** and **predictability** of the initially produced/perceived word and target correction, as well as the **phonemic** and **semantic** distance between these forms.

4.1 Speaker Model

A significant challenge in modeling whether or not a speaker has made correction in naturalistic contexts is that instances where a speaker *could* have made a correction but decided not to can only be reliably and accurately determined in highly controlled production studies. Here, we use comprehender corrections to approximate instances where a correction *could* have been made by the speaker. We make this assumption for two reasons. First, comprehender repairs give us a distribution over possible speaker repairs and *may* include the true repair. Furthermore, we assume that the speaker may be incentivized to correct when she believes there is a greater discrepancy between the speaker’s intended and comprehender’s inferred meaning respectively.

Consider the following corrections made by the original speaker and three of the comprehenders in the experiment. Based on observed corrections, we define a **critical window** at the position where either a speaker or a comprehender has made a correction.

3. *Speaker Corrected*: Well we also in this area seem to have a lot of retirees people who don’t

want the heat of Florida but don’t want the **cold** of the northeast

4. *Comprehender Corrected #1*: Well we also in this area seem to have a lot of retirees people who don’t want the heat of Florida but don’t want the **cold** of the northeast
5. *Comprehender Corrected #2*: Well we also in this area seem to have a lot of retirees people who don’t want the heat of Florida but don’t want the **chill** of the northeast
6. *Comprehender Corrected #3*: Well we also in this area seem to have a lot of **retired** people who don’t want the heat of Florida but don’t want the heat of the northeast

For each critical window, we compute the difference in frequency, predictability, phonemic, and semantic representations of the speaker’s initial production (e.g., *heat* or *retirees*) and the various comprehender responses in those positions (e.g., *chill* or *cold* in response to *heat* and *retired* in response to *retirees*). We develop the following maximally-converging mixed-effects logistic regression model (Barr et al., 2013) to predict correction decisions in speakers:

$$\text{SpeakerCorrected} \sim \Delta \text{LogProbability} + \Delta \text{Frequency} + \text{Semantic Distance} + \text{Phonemic Distance} + 1|\text{item} + 1|\text{critical window}$$

where Δ denotes the *difference* between the frequency and log probability of the speaker’s initial production and the comprehender’s response. We included random intercepts for item

and critical window to account for the participant-level variability at each correction site. We set `SpeakerCorrected = 1` for critical windows where the speaker did, in fact, make a repair (for example, 3–5). For all other instances, `SpeakerCorrected = 0` (for example, 6).

4.2 Comprehender Model

When modeling comprehenders’ corrections, we restrict the analysis to *speaker corrected* utterances. In particular, we only consider critical windows defined by the speaker’s repair (3–5), thus excluding responses such as (6). We compare the word in the critical window in comprehender’s final response (*chill* or *cold*) with the word in the perceived initial utterance (*heat*), and annotate it as *corrected* or *uncorrected*. We use the same metrics (relative frequency, predictability, semantic, and phonemic distance), which we calculate over the critical word in the initial utterance (*heat*) and the repair in speaker’s final production (*cold*), to predict whether a comprehender will make a correction in their final response.

We then use the following parallel model to predict correction decisions in comprehenders:

$$\text{ComprehenderCorrected} \sim \Delta\text{LogProbability} + \Delta\text{Frequency} + \text{Semantic Distance} + \text{Phonemic Distance} + 1|\text{item} + 1|\text{critical window} + 1|\text{subject}$$

4.3 Results

The results from the mixed-effects models reveal key asymmetries in how speakers and comprehenders decide to correct errors, with distinct influences from predictability, frequency, phonemic form, and semantic features. Model coefficients are summarized in Figure 9.

In the speaker model, the correction behavior is a function of the initially produced/perceived word and the comprehender’s final response. When the log probability difference between the comprehender’s final inferred response and their initial perceived word is higher, a speaker is less likely to self-correct ($\beta = -0.260, SE = 0.04, p < 0.01$). This suggests that if comprehenders could recover the intended meaning from the perceived erroneous form, the speaker is less likely to intervene. Additionally, greater semantic distance between the comprehender’s initial perception and their final response increases the likelihood of correction ($\beta = 0.227, SE = 0.08, p < 0.05$), indicating that speakers are more likely to correct when their

utterance creates a significant meaning deviation. Phonemic distance exhibits a similar effect, with greater phonemic dissimilarity between the comprehender’s perceived word and the corrected form leading to a higher probability of correction ($\beta = 0.270, SE = 0.05, p < 0.001$). Word frequency does not significantly impact speaker correction behavior ($\beta = -0.05, SE = 0.04, p = 0.21$), suggesting that lexical frequency alone is not a primary driver of corrective processes in production.

In contrast, the comprehenders’ corrections are based on the difference between the speaker’s erroneous production and their intended repair. Unlike speakers, comprehenders are more likely to correct when the speaker’s repair has a much higher log probability than the produced error, as indicated by a significant positive effect of log probability difference ($\beta = 0.278, SE = 0.07, p < 0.01$). This suggests that comprehenders rely more heavily on contextual expectations when detecting errors and overriding them. While semantic distance also increases the likelihood of correction ($\beta = 0.139, SE = 0.07, p < 0.05$), its effect is weaker compared to speakers, implying that comprehenders may be less sensitive to semantic deviations. Crucially, phonemic distance has the opposite effect in comprehension compared to production: while speakers are more likely to correct errors when phonemic distance is large, comprehenders exhibit reduced correction likelihood as phonemic distance increases ($\beta = -0.404, SE = 0.08, p < 0.001$), suggesting the comprehenders are more likely to correct errors that are phonemically similar to the intended meaning. As in the speaker model, frequency does not significantly influence comprehender corrections ($\beta = 0.0068, SE = 0.007, p = 0.92$).

5 Discussion

In this study, we develop *SPACER*, a parallel dataset of speech errors and repairs, which is designed to examine how speakers and comprehenders engage in error correction. We identified naturally occurring speech errors and repairs from speech corpus, and used web-based experiments to examine how speech errors are detected and corrected by comprehenders. We conducted exploratory analysis on the asymmetries of error correction strategies between production and comprehension. Specifically, we used linear models to predict whether lexical properties of speaker-produced




Figure 9: Model coefficients from speaker and comprehender logistic regression models. Error bars denote standard error. Significance: *** ($p < 0.001$), ** ($p < 0.01$), * ($p < 0.05$), ns ($p > 0.05$)

errors and repairs are predictive of how the errors would be corrected by comprehenders, and whether comprehender’s error correction could inform the speaker’s self-repair. Our analysis shows asymmetries between error correction by speakers and comprehenders.

Our dataset links production and comprehension and enables a principled comparison between the two modes. Previous studies have largely examined production and comprehension of errors separately (Ryskin et al., 2018; Dell, 1986; Levelt, 1983; Blank et al., 2018), making it difficult to assess whether and how these processes might share underlying principles. Our dataset bridges the gap by examining how speakers and comprehenders make choices about error correction on the *same* utterance. This parallel structure enables integrated analysis of correction strategies in production and in comprehension, providing a valuable resource for future studies.

Our exploratory analysis reveals key asymmetries in how speakers and comprehenders engage in error correction. Comprehenders are likely to correct errors that are phonemically similar to a more plausible alternative, or when the error is not supported by contextual cues. Speakers, on the other hand, are more likely to self-correct errors that might not be recoverable for comprehenders. The asymmetries might imply potential interplay between comprehension and production, arising from differing demands of the two modalities: while speakers correct their own speech to ensure commu-

nicative clarity given the message that they know they want to communicate, comprehenders may be able to use predictability and form-based cues to successfully recover the intended word. The inverse effects of phonemic distance in our models—where comprehenders correct more when errors are phonemically similar to an alternative, while speakers correct more when errors are phonemically distant from the intended meaning—suggest that interlocutors may engage in complementary error correction strategies. Future work can leverage *SPACER* for more principled computational models that simulate error correction as a rational inference process over various linguistic constraints and cues.

6 Limitations

The dataset is restricted to single-word substitution errors and does not include other common types of speech error such as insertions, deletions, and transpositions. While this allows for a controlled investigation of the error correction process, it may not capture other types of errors and correction strategies in naturalistic communication. Future work could expand *SPACER* to include a broader range of error types to better understand the full spectrum of production and comprehension repair mechanisms.

A key challenge when studying how speakers may adapt their correction strategies in naturalistic contexts is that the intended lexical target is often difficult to determine. In other words, while speech

error corpora provide *positive* examples of speaker corrections, they do not provide instances of *negative* examples i.e., where a speaker *could* have made a repair, but opted against it. Our corpus approach allows us to locate speech errors that have been overtly repaired, but we do not have access to speaker’s true communicative goal to identify uncorrected speech errors. We used the distribution of comprehender corrections as an approximation for speaker’s communicative goal, to identify errors that could have been corrected by the speaker but were left uncorrected. We acknowledge that the approximation may be different from the speaker’s intended meaning.

We used a web-based experiment to investigate how utterances are interpreted and corrected by comprehenders. There are several differences between our experiment and real-time language comprehension. First, comprehenders were presented with the full key sentence, whereas listeners do not have access to the entire utterance at once. Second, comprehenders are given unlimited time to make corrections. In contrast, real-time comprehension is constrained by limited cognitive and processing resources. Third, the text-based presentation method does not provide prosodic and phonetic cues. Comprehenders might rely on disfluencies and pauses to detect errors. Finally, it bears mentioning that the speaker and comprehender corrections are situated in entirely different contexts, separated by both time and space.

7 Acknowledgements

We thank the anonymous reviewers at CMCL, members of the QuantLang Lab at UC Irvine and TedLab at MIT, and the audience at Human Sentence Processing (HSP 2023), Society for Computation in Linguistics (SCiL 2023), and California Meeting on Psycholinguistics (CAMP 2024) for their valuable feedback on previous iterations of this work.

References

Jennifer E Arnold, Jason M Kahn, and Giulia C Pancani. 2012. Audience design affects acoustic reduction via production facilitation. *Psychonomic bulletin & review*, 19:505–512.

Matthew Aylett and Alice Turk. 2004. The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic

prominence, and duration in spontaneous speech. *Language and speech*, 47(1):31–56.

- Markus Bader and Michael Meng. 2018. The misinterpretation of noncanonical sentences revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(8):1286.
- Dale J. Barr, Roger Levy, Christoph Scheepers, and Harry J. Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3):255–278.
- Elizabeth R Blackmer and Janet L Mitton. 1991. Theories of monitoring and the timing of repairs in spontaneous speech. *Cognition*, 39(3):173–194.
- Helen Blank, Marlene Spangenberg, and Matthew H Davis. 2018. Neural prediction errors distinguish perception and misperception of speech. *Journal of Neuroscience*, 38(27):6076–6089.
- Laurel Brehm, Carrie N Jackson, and Karen L Miller. 2019. Speaker-specific processing of anomalous utterances. *Quarterly Journal of Experimental Psychology*, 72(4):764–778.
- Marc Brysbaert and Boris New. 2009. Moving beyond kučera and francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for american english. *Behavior research methods*, 41(4):977–990.
- Esteban Buz, Michael K Tanenhaus, and T Florian Jaeger. 2016. Dynamically adapted context-specific hyper-articulation: Feedback from interlocutors affects speakers’ subsequent pronunciations. *Journal of memory and language*, 89:68–86.
- Sasha Calhoun, Jean Carletta, Jason M. Brenier, Neil Mayo, Dan Jurafsky, Mark Steedman, and David Ian Beaver. 2010. The nxt-format switchboard corpus: a rich resource for investigating the syntax, semantics, pragmatics and prosody of dialogue. *Language Resources and Evaluation*, 44:387–419.
- Gary S Dell. 1986. A spreading-activation theory of retrieval in sentence production. *Psychological review*, 93(3):283.
- Jack Dempsey, Anna Tsiola, Suphasiree Chantavarin, Fernanda Ferreira, and Kiel Christianson. 2023. Nonce word evidence for the misinterpretation of implausible events. *Journal of Cognitive Psychology*, 35(5):526–544.
- Fernanda Ferreira, Karl GD Bailey, and Vittoria Ferraro. 2002. Good-enough representations in language comprehension. *Current directions in psychological science*, 11(1):11–15.
- Fernanda Ferreira and Janis Stacey. 2000. The misinterpretation of passive sentences. *Manuscript submitted for publication*. Publisher: Citeseer.

- Victor S Ferreira and Gary S Dell. 2000. Effect of ambiguity and lexical availability on syntactic and lexical production. *Cognitive psychology*, 40(4):296–340.
- Victoria Fromkin and Victoria Fromkin. 1973. *Speech errors as linguistic evidence*. Mouton The Hague.
- Richard Futrell. 2023. Information-theoretic principles in incremental language production. *Proceedings of the National Academy of Sciences*, 120(39):e2220593120.
- Richard Futrell and Edward Gibson. 2017. L2 processing as noisy channel language comprehension. *Bilingualism: Language and Cognition*, 20(4):683–684.
- Richard Futrell, Edward Gibson, and Roger P. Levy. 2020. Lossy-context surprisal: An information-theoretic model of memory effects in sentence processing. *Cognitive science*, 44(3):e12814.
- Hanna S Gauvin and Robert J Hartsuiker. 2020. Towards a new model of verbal monitoring. *Journal of Cognition*, 3(1).
- Edward Gibson, Leon Bergen, and Steven T. Piantadosi. 2013. Rational integration of noisy evidence and prior semantic expectations in sentence interpretation. *Proceedings of the National Academy of Sciences*, 110(20):8051–8056.
- John J Godfrey, Edward C Holliman, and Jane McDaniel. 1992. Switchboard: Telephone speech corpus for research and development. In *Acoustics, speech, and signal processing, IEEE international conference on*, volume 1, pages 517–520. IEEE Computer Society.
- Adele E Goldberg and Fernanda Ferreira. 2022. Good-enough language production. *Trends in Cognitive Sciences*, 26(4):300–311.
- Robert J Hartsuiker and Herman HJ Kolk. 2001. Error monitoring in speech production: A computational test of the perceptual loop theory. *Cognitive psychology*, 42(2):113–157.
- Gregory Hickok, John Houde, and Feng Rong. 2011. Sensorimotor integration in speech processing: computational basis and neural organization. *Neuron*, 69(3):407–422.
- T Jaeger and Roger Levy. 2006. Speakers optimize information density through syntactic reduction. *Advances in neural information processing systems*, 19.
- T Florian Jaeger. 2010. Redundancy and reduction: Speakers manage syntactic information density. *Cognitive psychology*, 61(1):23–62.
- Mark Koranda, Martin Zettersten, and Maryellen C MacDonald. 2018. Word frequency can affect what you choose to say. In *CogSci*.
- W Levelt. 1999. Producing spoken language. *The neurocognition of language*, pages 83–122.
- Willem JM Levelt. 1983. Monitoring and self-repair in speech. *Cognition*, 14(1):41–104.
- Roger Levy. 2008. A noisy-channel model of human sentence comprehension under uncertain input. In *Proceedings of the 2008 conference on empirical methods in natural language processing*, pages 234–243.
- Roger Levy, Klinton Bicknell, Tim Slattery, and Keith Rayner. 2009. Eye movement evidence that readers maintain and act on uncertainty about past linguistic input. *Proceedings of the National Academy of Sciences*, 106(50):21086–21090.
- Eric Meinhardt, Eric Baković, and Leon Bergen. 2020. Speakers enhance contextually confusable words. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 1991–2002.
- Kanishka Misra. 2022. minicons: Enabling flexible behavioral and representational analyses of transformer language models. *arXiv preprint arXiv:2203.13112*.
- Sieb Nooteboom and Hugo Quené. 2019. Temporal aspects of self-monitoring for speech errors. *Journal of Memory and Language*, 105:43–59.
- Sieb G Nooteboom and Hugo Quené. 2017. Self-monitoring for speech errors: Two-stage detection and repair with and without auditory feedback. *Journal of Memory and Language*, 95:19–35.
- Nazbanou Nozari, Gary S Dell, and Myrna F Schwartz. 2011. Is comprehension necessary for error detection? a conflict-based account of monitoring in speech production. *Cognitive psychology*, 63(1):1–33.
- Nazbanou Nozari, Michael Freund, Bonnie Breining, Brenda Rapp, and Barry Gordon. 2016. Cognitive control during selection and repair in word production. *Language, cognition and neuroscience*, 31(7):886–903.
- John K Pate and Sharon Goldwater. 2015. Talkers account for listener and channel characteristics to communicate efficiently. *Journal of Memory and Language*, 78:1–17.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Till Poppels and Roger Levy. 2016. Structure-sensitive noise inference: Comprehenders expect exchange errors. In *CogSci*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.

Ardi Roelofs. 2020. On (correctly representing) comprehension-based monitoring in speaking: Rejoinder to nozari (2020). *Journal of Cognition*, 3(1).

Rachel Ryskin, Richard Futrell, Swathi Kiran, and Edward Gibson. 2018. Comprehenders model the nature of noise in the environment. *Cognition*, 181:141–150.

Rachel Ryskin, Shukhan Ng, Katherine Mimnaugh, Sarah Brown-Schmidt, and Kara D Federmeier. 2020. Talker-specific predictions during language processing. *Language, cognition and neuroscience*, 35(6):797–812.

Yuhan Zhang, Rachel Ryskin, and Edward Gibson. 2023. A noisy-channel approach to depth-charge illusions. *Cognition*, 232:105346.

A Example items and illustration of the pre-processing pipeline

- Original utterance: it depends on whether **you** whether we figure that we have a defense oriented military or an **aggressive** aggression oriented military
- Frame 1: it depends on whether **you** whether we figure that we have a defense oriented military or an aggression oriented military
- Frame 2: it depends on whether we figure that we have a defense oriented military or an **aggressive** aggression oriented military

B Instructions and prompts provided to participants in the web-based editing experiment

Welcome!

We are annotating human speech data, and we need your help to find out typos or speech errors in the transcription.

In each trial, you will see some context, followed by a key sentence marked in **BOLD**.

There might be an error in some key sentences.

If you believe there is an error in the transcription, then you may correct it by replacing the erroneous word with a more appropriate choice.

Please do not add or delete any word, just swap the erroneous word with an appropriate replament.

If the sentence looks good, or you are unsure what is the intended meaning, leave it unchanged.

Most punctuation marks have been removed from the items that you will encounter (except periods and apostrophes).

Please do not add any new punctuation to, or change any capitalization in the target sentence, unless it is part of the word-level correction (e.g.when correcting I to you).

Figure 10: Instructions provided to participants at the beginning of the error correction experiment.

Speaker A: Uh you get a few of them that are a little bit sloppier uh just taking uh you know just junk items that they can just get you know ten to twenty bucks for real quick
 Speaker A: Probably the biggest items that uh that you get in our neighborhood or would be the high classed automobiles
 Speaker A: uh Mercedes Volvos things that they can get um you know big dollars for and jewelry are the major targets

Trial explained: Do not over-correct. Informal speech is filled with pauses or colloquial expression. Do not correct those filler words or information expression such as "big dollars". This sentence should be left unchanged, and you can move the confidence bar to continue to the next trial.

Speaker A: uh Mercedes Volvos things that they can get um you know big dollars for and jewelry are the major targets

Confidence: / 100

Figure 11: Sample practice trial with feedback to ensure familiarization with the above instructions.