Noisy-channel sentence comprehension in aphasia: the role of noise in the context

Rachel Ryskin (ryskin@mit.edu; MIT & BU), Edward Gibson (MIT), & Swathi Kiran (BU)

Recent psycholinguistic theories propose that communication is typically subject to noise (e.g., Ferreira & Patson, 2007; Gibson, Bergen, & Piantadosi, 2013; Levy, 2008)—due to speaker error, environmental noise, or listener misperception. Comprehenders rationally integrate noisy input (e.g., "The mother gave the candle the daughter.") with prior knowledge to infer the intended meaning of a sentence (e.g.,"The mother gave the candle *to* the daughter."). Critically, comprehenders have a model of the noise -- they expect some noise corruptions (e.g., deletion of a word) to be more likely than others (e.g., insertion of a word).

This noisy-channel framework has also been proposed to account for failures to use syntactic cues and an over-reliance on plausibility in persons with aphasia (PWA; Gibson, et al., 2015; Warren, et al. 2017). On this account, PWA have a higher base-rate of noise in their noise model. However, the underlying mechanism is not known. Perhaps PWA expect a higher rate of errors because of perceptual difficulties, or because they themselves produce errors in speaking, or perhaps they fail to update their noise model. In healthy adults, the noise model is tuned to the environment (Gibson et al., 2013; Ryskin et al., 2017). For example, readers make more noisy-channel inferences when there is more noise in the context (i.e., surrounding sentences contain errors). In order to understand how and why the noise model differs in aphasia, **we ask whether PWA can adapt their noise model to the local context**.

Methods: In 2 experiments (E1: N=40 online controls, E2: N=9 controls and N=16 PWA with filler accuracy > 60%, data collection ongoing), participants read 24 syntactically licensed but semantically implausible test sentences (e.g., The mother gave the candle the daughter) and answered comprehension questions (e.g., Did the daughter receive something?). Responses to these questions indexed whether participants were relying on the literal syntax (e.g., No: the candle received something) or their prior knowledge (e.g., Yes: daughters are more likely to receive something than candles), inferring that a more plausible intended sentence (The mother gave the candle to the daughter) had been corrupted by noise (e.g., deletion of "to"). Six test sentence types (from Gibson et al., 2013) were intended to elicit different plausible noise inferences (word deletion, insertion, or exchange). Test sentences were intermixed with exposure sentences which were error-free in the No Noise condition or contained errors (e.g., The fireman rescued the boy in the time of nick) in the Noise condition.

Results: Participants were less likely to make inferences (i.e., more likely to use the literal interpretation) on exchange sentences than deletion and insertion sentences (E1: deletion-exchange odds ratio[OR]=0.11; E2: OR=0.26; see Fig. 1 & 2). There was an overall effect of noise condition: participants in the Noise condition were less likely to interpret test sentences literally relative to the No Noise condition (E1: OR=0.43; E2: OR=0.61). Critically, PWA were less likely to interpret test sentences literally than controls (OR=0.38; i.e., more likely to rely on plausibility), and the effect of noise was reduced in PWA relative to controls (i.e., controls make more inferences in the Noise condition, OR=2.71, and PWA don't, OR=0.98).

Discussion: Participants in both experiments were more likely to make noisy-channel inferences when the noise corruption to be posited was a deletion than an insertion or exchange (replicating Gibson et al, 2013; 2015). A noisy-channel account of comprehension successfully explains patterns of comprehension in PWA (and controls). But, PWA make more inferences overall suggesting that they expect a higher base rate of errors in their input. Critically, unlike healthy controls, PWA appear not to modulate the rates of noisy-channel inferences they make based on the amount of noise in the input. A fruitful avenue for future research will be to determine whether the noise model is inaccurate in PWA because they fail to *update* their representation, or because they do not *perceive* errors.

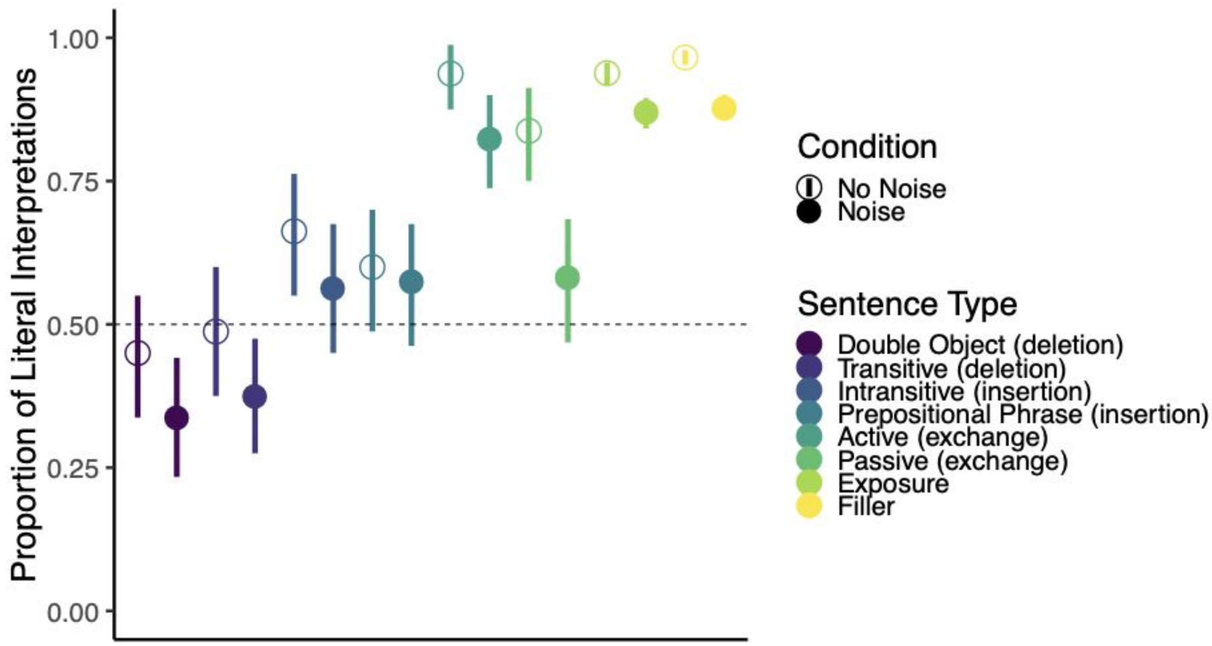

Figure 1. Experiment 1: Proportion of literal interpretations of test sentences by healthy controls on Amazon's Mechanical Turk, by sentence type (labeled with plausible noisy-channel correction in parentheses) and noise condition. Error bars indicate bootstrapped 95% confidence intervals around the mean.

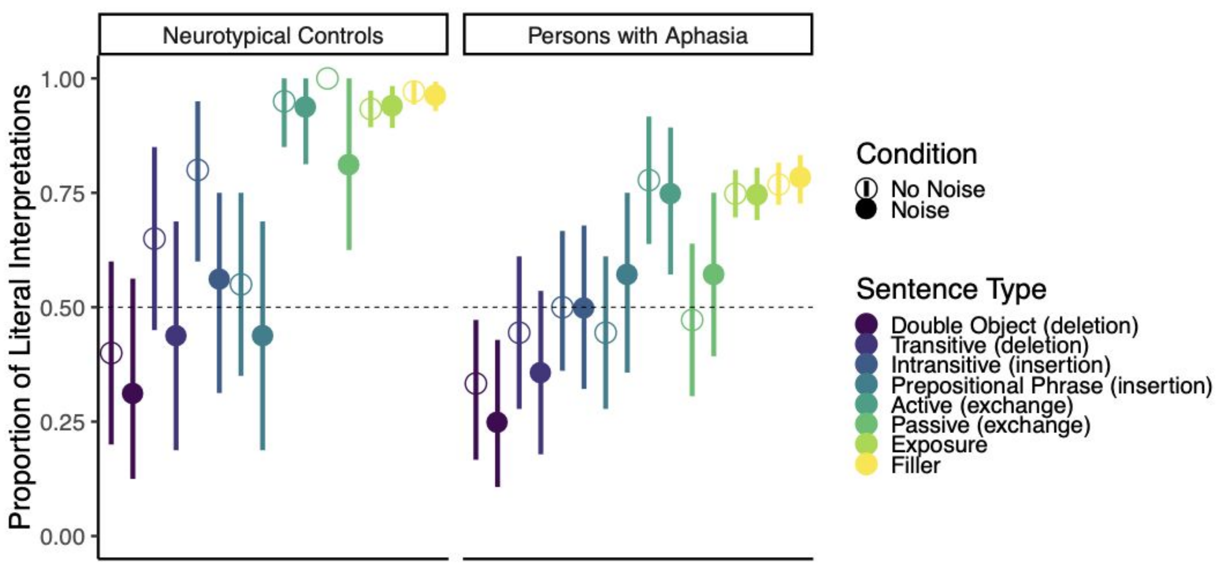

Figure 2. Experiment 2: Proportion of literal interpretations of test sentences by healthy controls (left) and persons with aphasia (right), by sentence type (labeled with plausible noisy-channel correction in parentheses) and noise condition. Error bars indicate bootstrapped 95% confidence intervals around the mean.