

Pronominal reference and inferred explanations: a Bayesian account

Background and Study Kehler & Rohde (2013) posit a Bayesian analysis of pronoun use whereby biases towards referents of pronouns ($P(\textit{referent}|\textit{pronoun})$) are determined by combining the prior probability that a referent will get mentioned next ('next-mention' biases; $P(\textit{referent})$) and the likelihood that a pronoun will be used to mention that referent ($P(\textit{pronoun}|\textit{referent})$). Crucially, the factors that condition these terms are different: Next-mention biases are determined primarily by semantically-driven factors (e.g., coherence relations), whereas the production bias is sensitive primarily to information structure and grammatical role (e.g., favoring pronominalizing mentions of subject v. other referents; Rohde 2008, Fukumura & van Gompel 2010, Rohde & Kehler 2013).

We examine the model using data from a passage completion task with a novel 2x2 relative clause (RC) x prompt type design (1a-b). The RC manipulation utilizes the fact that RCs attached to direct objects can be inferred to provide explanations of the matrix event (Rohde et al. 2011).

- (1) a. The boss fired the employee who was embezzling money. (He) _____
b. The boss fired the employee who was hired early last year. (He) _____

Although not entailed, (1a) invites the inference that the employee was fired because of the embezzling. Crucially, such inferences are unnecessary for sentence felicity (cf. 1b). The prompt either enforces pronominalization of the next-mentioned referent (pronoun prompt) or is unconstrained (free prompt).

Participants (n=40) were given a context sentence per (1a-b) along with gender-disambiguating pictures and asked to write a follow-on sentence (24 stimulus sets interleaved with 36 fillers). Target stimuli used object-biased implicit-causality (IC2) verbs with two same-gender referents. The continuations were annotated for coherence relation (explanation or other), next-mention (which event participant the continuation's matrix subject referred to, if either), and referential form (pronoun or other). Outcomes were modelled using mixed-effects logistic regression.

Hypotheses and Results Accounts that appeal primarily to surface-level characteristics of the context find little to distinguish (1a-b). The Bayesian analysis predicts a difference, however, based on an interconnected sequence of referential and coherence-driven interdependencies. First, it predicts that participants will write fewer explanation continuations in (1a) than (1b), since the RC in (1a) already provides a cause (Simner & Pickering 2005; Kehler et al. 2008; Bott & Solstad 2012). Second, this difference is predicted to yield a difference in next-mention biases: Since IC2 verbs impute causality to the object, a greater number of explanation continuations for (1b) should lead to a greater number of next-mentions of the object. Third, the RC manipulation is expected to affect which referent a pronoun refers to, since $P(\textit{referent}|\textit{pronoun})$ is determined in part by next-mention expectations. Fourth, pronoun production should not be similarly affected by the RC manipulation, but instead only by grammatical role. Finally, since pronoun interpretation is determined in part by production biases, more subject references are expected in the pronoun prompt vs. free condition. All of these predictions were confirmed.

Conclusion Pronoun interpretation biases are sensitive to the inference of implicit explanations whereas production biases are not, revealing precisely the asymmetry between interpretation and production predicted by the Bayesian analysis.