

Individual differences guide pronoun interpretation in semantically constraining contexts

Valerie J. Langlois & Jennifer E. Arnold (University of North Carolina – Chapel Hill)

valeriel@live.unc.edu

Previous research has established that syntactic and semantic cues influence the interpretation of third-person pronouns (Arnold 1998, 2001; Kehler et al., 2008; Stevenson et al., 1994). Some verbtypes exhibit a syntactic bias toward the subject. For example, in *Ana is cleaning up with Liz. She needs the broom*, listeners tend to assign “she” to Ana, the subject of the previous sentence. However, some verbs also impose a semantic bias. Following events with transfer verbs, there is a general subject bias, in addition to a relative preference to assign the pronoun to the goal character, as in Liz in *Ana threw the ball to Liz. She...* or Ana in *Ana got the ball from Liz. She...* (Stevenson et al., 1994). This raises a question: do people learn both of these biases in the same way, or are they fundamentally different?

Here we test this question by considering the role of experience with the kinds of references that are most frequent in natural language. Current models explain pronoun comprehension in terms of representations about who is likely to be mentioned next (Arnold et al., 2001; Kehler et al., 2008, Hartshorne et al., 2015). One idea is that people could learn that both subject and goal reference is more expected because both patterns are frequent (Arnold 1998, 2001). Support for this idea comes from evidence that for sentences like the “cleaning” example above, the subject bias is stronger for people with stronger print exposure, which is one domain of linguistic experience (Arnold et al., 2018). However, nothing is known about how print exposure affects pronoun comprehension following transfer events. Does it lead to a stronger subject bias, or a stronger goal bias? We test this question with both an offline interpretation experiment (Exp. 1) and a story completion task (Exp. 2).

Methods: In Exp1, 65 MTurk participants saw 60 trials, 24 of which included transfer-of-possession verbs (e.g. *threw, got*). Each trial presented a two-sentence story with images of the two characters and something associated with the event (Fig1). We manipulated verbttype, so that for half the trials the goal was the object-of-PP (e.g. *Ana threw the ball to Liz and then she fell down*), and for half it was the subject (*Ana got the ball ...*). The second sentence was always plausible for both characters, e.g. either character could have fallen down. After the story, participants answered two questions; the critical one measured pronoun comprehension (e.g. Who fell down?). In Exp2, the same two-story sentences were presented visually to the participant, but were cut off right after the pronoun; 70 MTurk participants were asked to type out a plausible continuation. 12% of the data was excluded from analysis due to ambiguous interpretation of the continuation. The remaining trials were coded as either a goal continuation or a source continuation. In both experiments, print exposure was measured by the Author Recognition Task (ART; Stanovich & West, 1989). Participants saw a list of real and fake authors, and were asked to indicate which ones they recognized as real. ART scores were calculated based on the # of correct minus the # of incorrect.

Results: In Exp1, there was a goal bias, in that participants interpreted the pronoun as the goal (56%) over the source (44%; Fig2a), as well as an overall subject bias. People with higher ART had a higher the rate of selecting the subject regardless of thematic role ($p=.025$; Fig. 2c). In Exp. 2, participants continued talking about the goal (77%) more than the source referent (23%; Fig3a). In addition, the rate of mentioning the subject character (regardless of thematic role) was higher for participants with higher ART ($p= .012$; Fig3c).

Conclusions: We found that linguistic experience correlates with use of the subject bias, but not the goal bias. This suggests that the subject bias may be learned from exposure to language, but not the goal bias. Findings broadly support the importance of referential frequency for the learning of discourse conventions like the subject bias. This is consistent with models that use referential probability to explain pronoun comprehension. However, judgments about semantic role predictability may be guided not by frequency, but by on-the-fly judgments about which continuation is more semantically likely (Hartshorne et al., 2015).

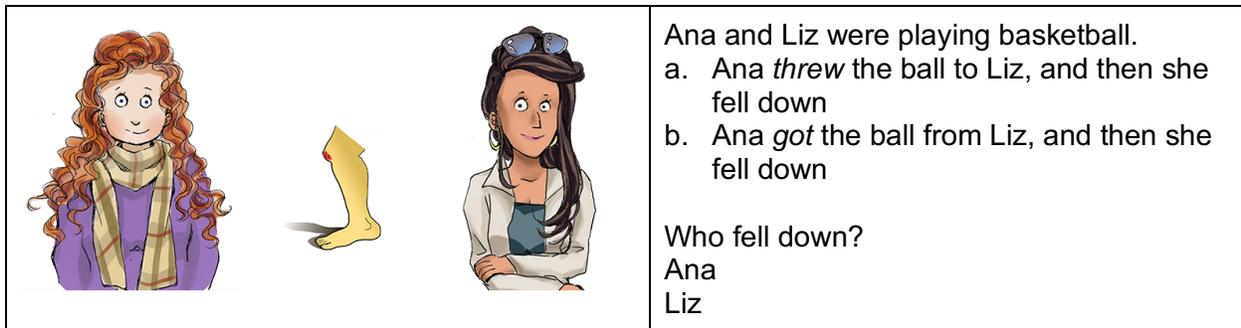


Figure 1: Example of a critical trial. Left: Example visual stimulus. Right: Examples of the story participants heard and the comprehension question.

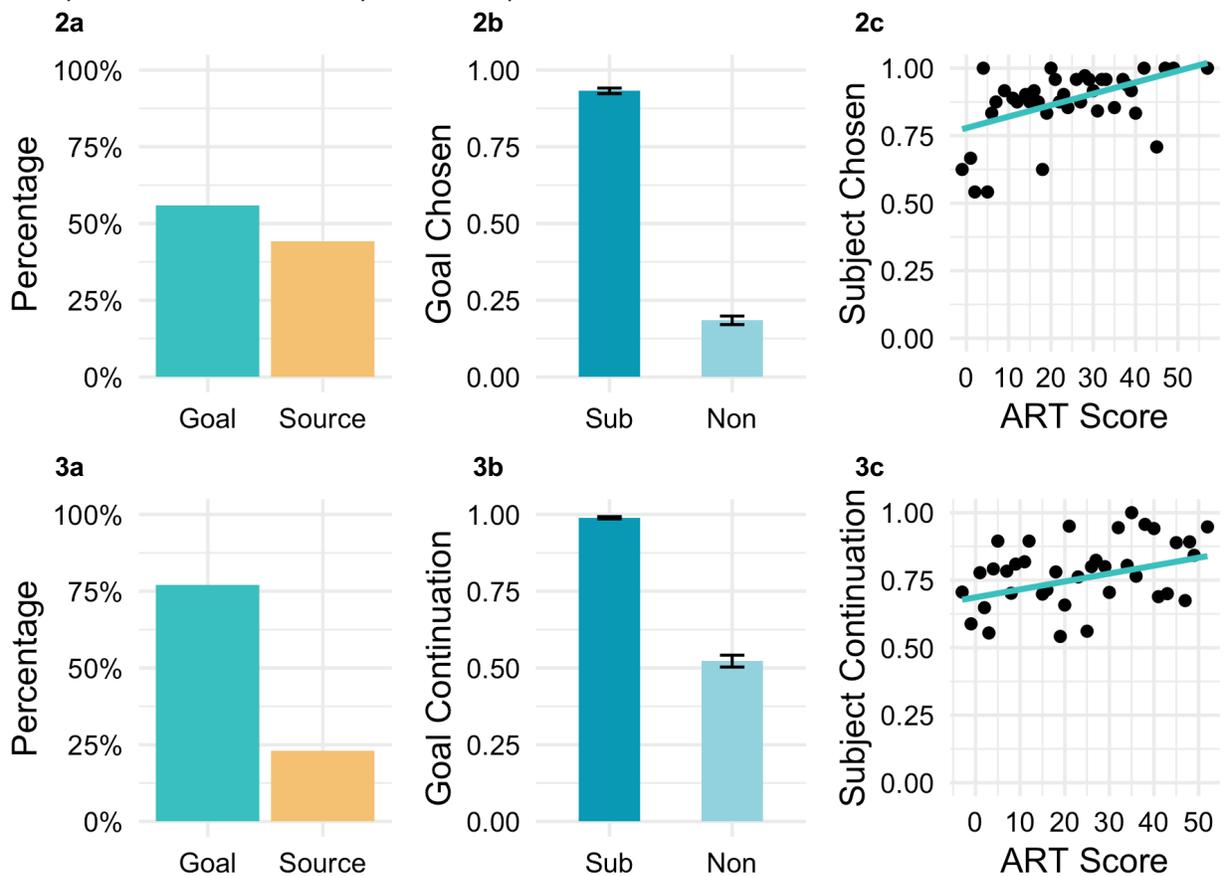


Figure 2 (Exp. 1 results) and Figure 3 (Exp. 2 results). a. Percentages of goal referents and source referents chosen. b. Breakdown of goal referents chosen between subject position and non-subject position. Error bars represent standard errors around the mean. c. ART score predicts subject bias

References:

Arnold, J. E. (1998). Reference Form and Discourse Patterns. Dissertation, Stanford University.
 Arnold, J. E. (2001). The effects of thematic roles... *Discourse Processes*, 31(2), 137-162.
 Arnold, J. E. (2018). Linguistic experience affects pronoun interpretation. *JML*, 102, 41-54.
 Hartshorne, J. K. et al. (2015). The causes and consequences *LCN*, 30, 716-734.
 Kehler, A. et al. (2008). Coherence and coreference revisited. *Journal of semantics*, 25(1), 1-44.
 Stanovich, K. E., & West, R. F. (1989). Exposure to print *Reading Res. Quart.*, 24, 402-433.
 Stevenson, R. J. et al. (1994). Thematic roles, focus *LCP*, 9, 519-548.