

The role of definiteness in ad hoc implicatures

Summary: This study investigates how ad-hoc implicatures and the definiteness presupposition of *'the'* interact. Using a truth value judgment task (Crain & Thornton 2000), we examine whether English-speaking adults interpret “*Mary bought a striped sweater*” differently from “*Mary bought the striped sweater*” in a context where there are two possible referents, one which is best described with one adjective (e.g. “*striped*”) and the other which is best described with two adjectives (e.g. “*striped and spotted*”). Contrary to what standard models of implicature generation predict, we find that uses of *'the'* are rejected more frequently than uses of *'a'* when the item bought is best described with two adjectives. This shows that the use of the indefinite blocks the generation of potential ad-hoc implicatures, which suggests that the processing of presuppositional content takes precedence over the processing of (ad-hoc) implicatures.

Ad-hoc implicatures and reference disambiguation: Under standard accounts of meaning enrichment, ad-hoc implicatures (Hirschberg 1991) are generated by treating a contextually provided alternative as false. When p is used in a context where $p \wedge q$ is a relevant alternative, an implicature that $\neg(p \wedge q)$ is generated. For example, in a context where there is a person with glasses and a person with both glasses and a hat, adults and even preschool-aged children interpret “*My friend has glasses*” as referring to the person with only glasses (Stiller et al. 2015). The ad-hoc implicature appears to disambiguate the two possible referents who both match the literal interpretation of “*My friend has glasses.*” Note, however, that the denotation of “*my friend*” independently implies there is a unique relevant friend being described.

Manipulating uniqueness: In this study we test whether ad-hoc implicatures provide reference disambiguation when the description of the possible referent doesn't imply uniqueness. We compare how the definite article, *'the'*, and the indefinite article, *'a'*, are interpreted when two contextually provided referents match the literal denotation of the NP. In a scenario where there is a sweater with stripes and a sweater with both stripes and spots (Figure 1), we assess how adult English speakers interpret (1) and (2):

- (1) Mary bought **the** striped sweater.
- (2) Mary bought **a** striped sweater.




Figure 1. Critical target image, paired with (1) in the definite condition and (2) in the indefinite condition.

Experiment: Participants: 60 English native speakers were recruited through Prolific and randomly assigned to either the *'a'* or *'the'* condition. Participants were paid at an average pay rate of £11.75/hour for the task, which took on average 6m9s to complete. Procedure: The task was a truth value judgment task, implemented and hosted on Qualtrics. Participants were given a back story about characters who were shopping at the store. On each trial, they saw a picture containing three items, and a shopping basket under one of the items. A puppet named Raffie described which item the character purchased (using either a definite or an indefinite description), and participants had to indicate whether Raffie was right or wrong by clicking on ‘Yes’ or ‘No’. Materials: Noun phrase type (definite *'the'* vs. indefinite *'a'*) was a between-subject variable. Critical target trials involved weak/under-informative descriptions containing one adjective, such as “*Mary bought a/the striped sweater*”, to describe a context in which there was both a sweater with stripes and a sweater with stripes and spots, and Mary had bought the one with stripes *and* spots (see Figure 1, paired with (1)/(2)). If participants computed the ad-hoc implicature that the sweater Mary bought didn't contain spots, they were expected to reject the test sentence; if not, they would accept the test sentence on its literal meaning. The experiment also included

unambiguously true and unambiguously false 1- and 2-adjective controls, in which the test sentences were clearly true or clearly false descriptions of the purchased item (see Figures 2 and 3 for examples). We also included clearly true and clearly false filler items which involved descriptions that did not contain any adjectives. In all, each participant saw 2 training items, followed by 30 test items: 12 ambiguous target trials containing either 'a' or 'the', 6 clearly true/clearly false 1-adjective controls, 6 clearly true/clearly false 2-adjective controls, and 6 adjective-less fillers. Order of presentation was randomized across participants.




Figure 2. Clearly true control image paired with the sentence: 'Max bought a/the plain shirt'.

Results: One participant was excluded for failing to score at least 12/18 (two thirds) accuracy on the unambiguous control and filler trials, leaving a total of 59 participants for analysis (29 in the 'a' condition and 50 in the 'the' condition). For these participants, accuracy was above 95% for all unambiguous filler and control conditions. Figure 4 displays the average proportion of yes-responses in the target 'a' and 'the' conditions (dots represent individual participant means). Mean acceptance in the indefinite 'a' condition was 93%, compared with 55% in the definite 'the' condition. We fitted a mixed effect logistic regression model on responses to the target conditions, with definiteness as a fixed effect, and random effects for subject and item. Model comparisons revealed a significant effect of definiteness ($\chi^2(1)=15, p<.0001$), with participants more likely to reject the underinformative target statements when it contained the definite article 'the'.




Figure 3. Clearly false control image paired with the sentence: 'Ellie bought a/the rainbow-coloured and polka-dotted dress.'

Discussion: The indefinite article was more referentially ambiguous than the definite article. Participants accepted 'a' more often than 'the' in scenarios where two referents matched the literal (at-issue) description of the purchased object. Standard accounts of implicatures predict no difference in how (1) and (2) disambiguate the object NP. Both have a contextually provided alternative, "the/a striped and spotted sweater", and negating these alternatives should pick out the sweater with stripes and no spots. Our findings suggest that computing presuppositional content blocks the generation of the ad-hoc implicature that would disambiguate the referent.




Figure 4. Performance on critical 'a' and 'the' targets.

This could be because 'a' carries a non-uniqueness presupposition (Hawkins 1978), or because listeners are sensitive to why terms with stronger presuppositions are avoided (Heim 1991). Either way, presuppositional processing takes precedence over implicature generation.

References: Crain, S., & Thornton, R. (1998). *Investigations in universal grammar: a guide to experiments on the acquisition of syntax and semantics*. (Language, speech, and communication).
 • Hawkins J. (1978). Definiteness and indefiniteness: a study in reference and grammaticality prediction
 • Heim, I. (1991). Artikel und Definitheit [Articles and definiteness]. In A. von Stechow & D. Wunderlich (eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*.
 • Hirschberg, J.L. (1991). *A theory of scalar implicature*.
 • Stiller, A.J., N.D. Goodman & M.C. Frank. (2015). Ad-hoc implicature in preschool children. *Language Learning and Development* 11.