

Probabilistic integration of linguistic framing in ad-hoc pragmatic implicatures

Authors: Andrés Gómez Emilsson, Michael C. Frank, Noah Goodman, Christopher Potts

Language comprehenders make a wide variety of pragmatic inferences to determine the meaning of ambiguous utterances in context. These inferences are affected by both the informativeness of the alternative messages with respect to an interpretation and the relevance of potential interpretations within the context (Grice, 1975). We use simple referential games to explore the interaction of these two factors, investigating a previously-studied informativeness inference similar to a scalar implicature (Stiller, Goodman, & Frank, 2011).

Participants in our experiments see a set of target items, e.g. three faces, one with no features, one with a hat, and one with a hat and glasses. They must guess which referent is the target of a description (e.g. “glasses”). To avoid confounds related to idiosyncratic preferences, we randomly varied the base item, named property, and relative position of items. A Bayesian model allows for quantitative predictions in these games by combining empirically-determined relevance values with model-derived informativeness in context (Frank & Goodman, 2012).

In **Experiment 1**, we told participants that “Bob can only say one word to communicate with you and he says: X” where X was the target feature. The goal for this experiment was to measure whether participants would infer an ad-hoc informativeness scale, leading to a scalar inference (“glasses” means *glasses and no hat*). 83% of participants from Amazon Mechanical Turk (N=322) chose the implicature-consistent referent, replicating previous work.

In **Experiment 2**, we tested alternative linguistic framings for this inference. Bob either says: “My *favorite* friend has glasses” or “My *least favorite* friend has glasses.” 53% (106/200) and 80% (77/96) of the participants chose the implicature-consistent option, respectively. We hypothesized that this manipulation affected participants’ *a priori* estimate of which referents were more likely to be chosen (the relevance distribution used by Frank & Goodman, 2012).

In **Experiment 3** we tested this hypothesis by asking what participants believe about Bob's preferences, i.e. which of the items they thought was Bob's favorite or Bob's least favorite. In the *favorite* case, 65% (44/68), 22% (15/68) and 13% (9/68) subjects chose the two-feature, one-feature and no-feature items, respectively. And for the *least favorite* case, the results were 21% (14/66), 14% (9/66) and 65% (43/66), respectively.

In the different framings of Experiment 2, participants appeared to be incorporating a bias that favors referents with either more (“favorite”) or fewer (“least favorite”) items. The Bayesian model captured a substantial part of this trend, with predicted values of 50% and 67% (vs. empirical values of 53% and 80%). These data suggest that pragmatic inferences reflect a graded integration of information about contextual relevance and linguistic informativeness.