

Mentioning atypical properties of objects is communicatively efficient

Elisa Kreiss¹, Judith Degen², Robert X.D. Hawkins², and Noah Goodman²

¹University of Osnabrück, ²Stanford University

What governs how much information speakers include in referring expressions? One pressure is for speakers to include just enough information for the interlocutor to correctly select an intended referent from among a set of potential referents [3]. This amounts to calling the target object a “banana” in 1a), where there is no competing banana; but a “yellow banana” in 1c), where there is a competing (brown) banana. However, speakers also have a well-documented preference to mention properties of objects “overinformatively”, especially color [4]. For example, speakers are likely to call the banana in 1a) a “brown banana” some of the time. More precisely, speakers tend to mention atypical rather than typical properties of objects overinformatively [5] [6] [7].

An account of why more typical properties are less likely to be mentioned is lacking. Some have proposed that it is due to a speaker-internal pressure to mention salient properties; others have proposed that speakers aim to facilitate the listener’s visual search. We ask: when should a rational speaker with the goal of communicating an intended referent mention an object’s color?

Model. We model reference production within the Rational Speech Act framework [1]. Taking inspiration from [2], utterances (simple nouns like “banana”, simple color adjectives like “blue”, and modified noun phrases like “blue banana”) are taken to have a graded semantics: rather than assuming that the bananas shown in Fig. 1a)-1d) are equally good instances of “banana” or that all shades are equally “blue”, we empirically elicited object-utterance typicality values on MTurk for all possible utterances. The pragmatic speaker selects utterances proportionally to the probability that a literal listener using a graded semantics will select the correct object. The listener is more likely to select a typical yellow banana upon hearing “banana,” thus it is more informative for the speaker to mention “COLOR banana” when the intended referent is atypical.

Production experiment. In order to evaluate the RSA model quantitatively, we collected freely produced referring expressions in a multi-player online reference game experiment using contexts such as those depicted in Fig. 1. 60 pairs of participants were recruited through MTurk and randomly assigned to speaker and listener role. Speakers used a chat window to produce a referring expression that would allow the listener to click on the target object. Once listeners made a choice by clicking on an object, feedback was provided to both participants. Stimuli were photo-realistic depictions of food items that occurred in three different colors, which differed in typicality. Conditions differed in whether mentioning color was “informative” (necessary for uniquely establishing reference, 1c-d) or “overinformative” (1a-b); and whether there was a competitor from another food category of the same color (1b/1d) or not (1a/1c).

Results are visualized in Fig. 2. For ease of exposition, we focus on whether or not color was mentioned at all (though the RSA model predicts the entire utterance distribution for each of the unique 1085 contexts). Color was mentioned more often in informative than in overinformative contexts ($\beta=5.27, p<.0001$) and more often when there was no color competitor than when there was ($\beta=.67, p<.0001$). Crucially, there was a main effect of typicality in the expected direction – the more typical an object was for the simple nominal expression, the less likely color was mentioned ($\beta=-4.11, p<.0001$), replicating previous studies. This was the case even when color was informative – in these cases, participants preferred to sometimes say “banana” for the very typical banana even though there was another banana present. BDA suggests the graded semantics model captures these data much better than a deterministic semantics model ($r=.8$).

We conclude that the systematicity with which speakers redundantly mention color implicates a system geared towards communicative efficiency rather than towards wasteful

overinformativeness. We discuss potential extensions of this approach to other production phenomena, such as optional instrument mention.


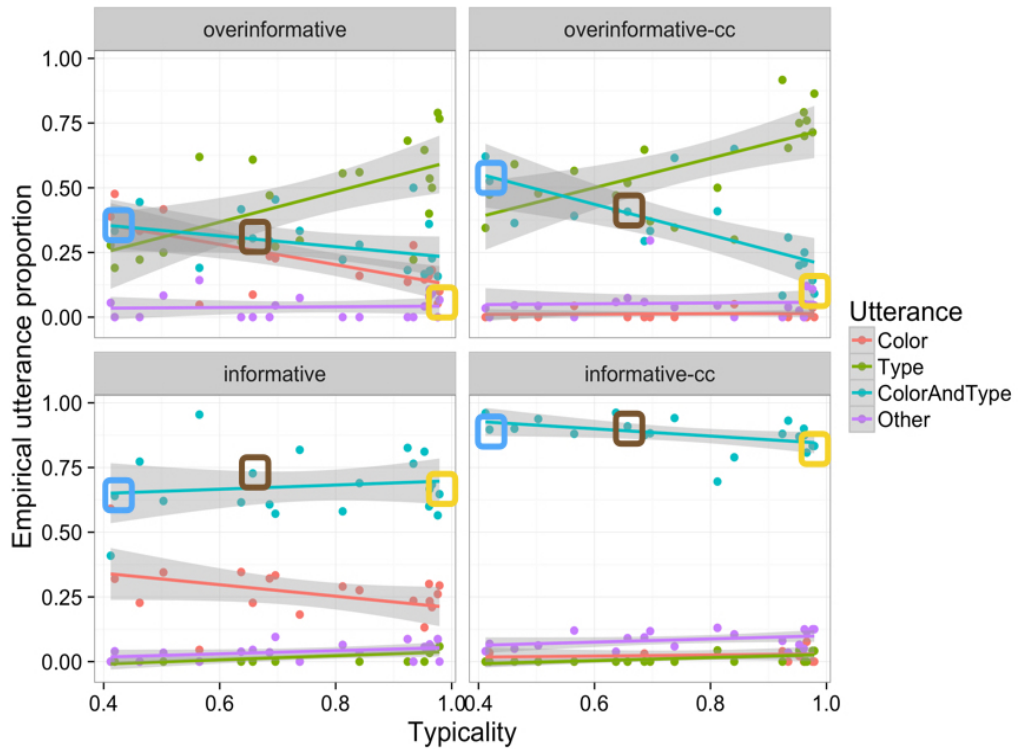


Fig. 1. Examples of relevant informativeness and color competitor presence conditions.

Fig. 2. Proportion of Color ("blue"), Type ("banana"), and ColorAndType ("blue banana") utterances as a function of mean object typicality for the Type utterance, across conditions. "COLOR banana" cases are circled in their respective color.



- [1] Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336, 998.
- [2] Graf, C., Degen, J., Hawkins, R. X. D., & Goodman, N. D. (2016). Animal, dog, or dalmatian? Level of abstraction in nominal referring expressions. In *Proceedings of CogSci 38*.
- [3] Grice, H. P. (1975). Logic and Conversation. *Syntax and Semantics*, 3, 41–58.
- [4] Pechmann, T. (1989). Incremental speech production and referential overspecification. *Linguistics*, 27(1), 89–110.
- [5] Rubio-Fernandez, P. (2016). How redundant are redundant color adjectives? An efficiency-based analysis of color overspecification. *Frontiers in Psychology*, 7 (153).
- [6] Sedivy, J. C. (2003). Pragmatic versus form-based accounts of referential contrast: evidence for effects of informativity expectations. *Journal of psycholinguistic research*, 32(1), 3–23.
- [7] Westerbeek, H., Koolen, R., & Maes, A. (2015). Stored object knowledge and the production of referring expressions: the case of color typicality. *Frontiers in Psychology*, 6(July), 1–12.