

Greta is a female director: When gender stereotypes interact with informativity expectations

Hailin Hao (hailinha@usc.edu)

Department of Linguistics, University of Southern California
Los Angeles, CA 90089-1693 USA

Muxuan He (muxuanhe@usc.edu)

Department of Linguistics, University of Southern California
Los Angeles, CA 90089-1693 USA

Zuzanna Fuchs (zfuchs@usc.edu)

Department of Linguistics, University of Southern California
Los Angeles, CA 90089-1693 USA

Abstract

Our world knowledge is deeply embedded in language. The word “banana” immediately brings to mind the color “yellow”. Oftentimes, however, when we read about “banana”, we do not expect the color “yellow” to be explicitly mentioned as in “yellow banana”. Frequent omission of obvious information given our world knowledge is predicted by so-called *informativity expectations*, which capture our preferences for newsworthy and informative utterances. Here we present two studies investigating informativity expectations in a socially situated context of gender stereotypicality and testing whether expressions like “female nurses” are less preferred than expressions like “female directors”. While the results show a clear impact of informativity expectations, effects of gender stereotypes turn out to be difficult to overcome.

Keywords: informativity; gender stereotypes; pragmatics; psycholinguistics; social psychology

Introduction

Comprehenders make predictions all the time when they process language. For instance, one would predict that an incomplete sentence like “This morning I had coffee with ...” is more likely to be followed by “milk” rather than “vinegar”, based on world knowledge (e.g., McRae & Matsuki, 2009). However, predictions during language comprehension are not guided solely based on world knowledge. Comprehenders may also expect that which is more informative and newsworthy (Rohde, Futrell, & Lucas, 2021; Brown & Dell, 1987; Grice, 1975, etc.). For example, the sentence “Emma bought a few yellow ...” might be expected to be followed by “notebooks” more than “bananas”. Given that bananas are prototypically yellow, the explicit mention of the prototypical color is oftentimes considered redundant and trivial (Sedivy, 2003; c.f. Rubio-Fernandez, 2019). The current study aims to test so-called *informativity expectations* (Rohde & Rubio-Fernandez, 2022), whereby comprehenders expect more informative and newsworthy use of language. In the two language experiments we present, we examine whether informativity expectations are modulated by occupational gender stereotypes in Mandarin Chinese.

Informativity Expectations

Informativity expectations, as an important component of language processing, have been extensively investigated in referent communication and identification tasks (Rubio-Fernandez, 2016, 2019; Engelhardt, Bailey, & Ferreira, 2006; Sedivy, 2003; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995, etc.). In a referent identification task, participants are usually given a visual context wherein a few objects are listed that have color and size similarities or differences. If an object is unique in color or size features among all the objects in the visual context, describing its color or size is considered redundant use of modification. Prior work has shown that both redundant and insufficient descriptions are less acceptable in rating tasks (e.g. Engelhardt et al., 2006) or that they slow referent identification, e.g., in eye-tracking experiments (Sedivy, 2003). In a more recent study, Rohde and Rubio-Fernandez (2022) examined how informativity expectations guided comprehenders’ predictions following a color term, and the results revealed that food items (which have prototypical colors) were more frequently chosen in a sentence-completion task and more looked at in an eye-tracking study, compared to clothing items (which have non-prototypical colors).

Occupational Gender Stereotype Effects in Language Processing

Occupational gender stereotypes are often used as a testing ground for the role of social factors in language processing. For example, von der Malsburg, Poppels and Levy (2020) ran a large-scale study investigating whether participants’ belief of the gender of the next government leader would affect their choice and comprehension of a *she* pronoun that co-refers with *president* during the 2016 United States presidential election and the 2017 United Kingdom general election, both of which had a popular female candidate. Results from the U.S. show that the pronoun *she* was rarely chosen in their sentence-completion task and that it increased processing difficulty in a self-paced reading (SPR) task. Results from

UK show that while the pronoun *she* was a preferred choice in sentence-completion, it was still more difficult to process in SPR. The findings suggest that long-term gender stereotypes showed more reliable effects. Papineau, Podesva and Degen (2022) analyzed how individual differences in gender ideologies and political affiliations influenced the use and processing of gender-neutral titles (e.g., *congressperson*). Although they did not find any effects of individual differences in online processing in the SPR task, there was evidence that a participant’s political views align with their choice of gender-neutral titles in a sentence-completion task. These findings indicate that gender stereotypes were inextricably connected with language; when we use and process language, it reflects our gender-related stereotypes.

Current Study

The current study asks whether the effects of informativity expectations can be found in gender stereotypes. Extending prior studies on informativity, our study aims to investigate the effects of informativity in a social context. Compared to gender stereotype effects observed in prior language processing research, our study examines a less-studied phenomenon in a non-English language. In the case of gender and occupation, informativity expectations may encourage predictions of gender-astereotypical occupations when a gender modifier is present, so that the occupation “director” might be more expected following “Greta is a female ...” vs. “Greta is a ...”. We investigate how the presence or absence of a gender modifier influences prediction of the occupation word that immediately follows by leveraging the optional gender modifier in Mandarin Chinese (henceforth, Mandarin). Examples in (1) illustrate how the gender modifier is used in Mandarin and its optionality. The gender modifier (i.e., *nv* ‘female’ or *nan* ‘male’) must match the gender of the proper name (i.e. *Mary* or *John*), but whether to use it or not is optional, regardless of whether the occupation is stereotypically associated with a particular gender.

- (1) a. *Mary shi yi-ge (nv) feixingyuan.*
 Mary is one-CL (female) pilot
 Mary is a (female) pilot.
 b. *John shi yi-ge (nan) baomu.*
 John is one-CL (male) babysitter
 John is a (male) babysitter.

In Chinese society, *feixingyuan* ‘pilot’ is a male-dominated occupation while *baomu* ‘babysitter’ is a female-dominated occupation. It would thus be regarded as economic ~~rational~~ use of language to use the gender modifier in these uncommon cases where a female takes on a male-dominated occupation (1a) and a male takes on a female-dominated occupation (1b).

To systematically examine the role of the gender modifier in the prediction of occupations, we manipulated the gender modifier by varying whether it is present or absent before the occupation. We tested two hypotheses regarding the role of

informativity, adapting Rohde and Rubio-Fernandez’s (2022) terminology. The Stereotypical Knowledge Hypothesis posits that knowing the gender of a person activates gender-stereotypical concepts including occupation. This hypothesis suggests that, with or without a gender modifier, (fe)male-stereotypical occupations are more predictable following (fe)male-biased names. However, the Informativity Hypothesis posits that the presence of a gender modifier motivates reasoning about why the modifier is used. This second hypothesis suggests that gender-stereotypical predictions will be lessened when a gender modifier is present vs. when it is absent.

We assess whether participants show informativity expectations using a forced-choice sentence completion task (Experiment 1) and a real-time sentence comprehension experiment (Experiment 2). In the **sentence completion task**, we ask participants to choose between a female-stereotypical occupation and a male-stereotypical occupation to complete a sentence. The Stereotypical Knowledge Hypothesis predicts no difference between stereotypical occupation responses across conditions. The Informativity Hypothesis predicts fewer stereotypical occupation responses when the gender modifier is present vs. when it is absent. The **reading experiment** investigates the informativity expectations with real-time processing data. If participants prefer the non-redundant use of the gender modifier, as the Informativity Hypothesis predicts, longer reading times (RTs) are expected at the gender-stereotypical occupation nouns vs. at the gender-astereotypical occupation nouns following the gender modifier. If the presence of gender modifier does not make a difference, as the Informativity Hypothesis predicts, similar RTs are expected across conditions.

Experiment 1: Sentence-Completion Task

Method

Participants We recruited 104 participants (76 female, 28 male) online. All participants self-identified as native speakers of Mandarin Chinese and as currently residing in Mainland China.

Table 1. Example target items in Experiment 1

	GENDERname	MODIFIER
a.	female	absent
	Wang Hong shi yi-ge ____ . (“Wang Hong is a ...”)	
b.	female	present
	Wang Hong shi yi-ge nv ____ . (“Wang Hong is a female ...”)	
c.	male	absent
	Wang Wei shi yi-ge ____ . (“Wang Wei is a ...”)	
d.	male	present
	Wang Wei shi yi-ge nan ____ . (“Wang Wei is a male ...”)	

Design & Materials Experiment 1 is a forced-choice sentence completion task. We adopted a within-participant 2x2 design crossing Gender (i.e., the gender bias of the protagonist’s name; female-biased vs. male-biased) and Modifier (i.e., the presence or absence of the gender modifier; present vs. absent). When the gender modifier was present, it was consistent with the gender indicated by the name (see Table 1).

For each target trial (illustrated in Figure 1), participants saw the sentence written out in Chinese characters, then there was a blank at the end of the sentence and participants were asked to complete the sentence by choosing between two options: a female-stereotypical occupation and a male-stereotypical occupation. All names and occupations were pre-normed (N = 40) to ensure the intended gender expectations.




Figure 1. Forced-choice sentence completion task (English translation included for the purpose of presentation)

We created 24 target items in the format illustrated in Table 1 and 48 fillers. We created 4 lists following a Latin square design to ensure that each participant saw each item in only one condition. Half of the target items used female-biased names and the other half used male-biased names. 12 of the fillers were attention check questions that target various kinds of general knowledge. 12 of the fillers were eye-catching celebrity-related questions. 24 of the fillers were concerned with pronoun resolution, half unambiguous and half ambiguous. We included ambiguous trials so that participants would not think that our task always required a single correct answer.

Procedure The experiment was hosted and administered online through Qualtrics. Each participant was randomly assigned one of the four lists. Item presentation was randomized and used a single-question-per-page format. The (fe)male-stereotypical occupation was randomly sampled from a list of pre-normed (fe)male-stereotypical occupations. The order of the two choices on the screen was counterbalanced.

At the start of the experiment, participants were given instructions for the sentence-completion task and completed practice trials. Participants were told that they should base their choices on their own preferences. To progress through the trials, participants chose a continuation and then clicked the “next” button to finalize their choice and proceed to the next trial. Final responses to the trials were recorded. On average, the experiment took about 8 minutes.

Analysis

Exclusion Three participants failed to answer correctly at least 80% of the attention check questions and were removed from all analyses.

Statistical Model We analyzed the binary responses of female-stereotypical versus male-stereotypical occupation, using logistic mixed-effects regression models. We included Gender and Modifier as fixed effects, and random intercepts for participants and items. For random slopes, we started from the maximal linear hierarchical model and then reduced it via model comparison. We report the maximal model that allowed for convergence (Barr, Levy, Scheepers, & Tily, 2013). All analyses were carried out in R using *lme4* (Bates, Mächler, Bolker, & Walker, 2015). In all analyses, significance was determined using Satterthwaite’s method of approximating degrees of freedom in the *lmerTest* package (Kuznetsova, Brockhoff, & Christensen, 2017). The significance level was set at .05.

Results

Results of the binary choice responses are plotted in Figure 2. The y-axis represents the proportion of female stereotypical occupation responses. Visual inspection of the data shows that (1) female names induce a much higher proportion of female-stereotypical occupation responses than male names, and (2) the presence of the gender modifier increases the proportions of *astereotypical* occupation responses. The results of the statistical analyses are consistent with these observations. Results revealed a main effect of Gender ($\beta = -2.510$, $SE = 0.247$, $z = -10.153$, $p < .001$): Overall, participants favored the occupation responses that matched the stereotypical occupation of the gender associated with the proper name. We did not find a main effect of Modifier ($\beta = 0.169$, $SE = 0.100$, $z = 1.683$, $p = 0.092$), but we found a significant interaction between Gender and Modifier ($\beta = 1.455$, $SE = 0.214$, $z = 6.797$, $p < .001$): Participants showed reduced preference for the gender-stereotypical occupations when the gender modifier was present.

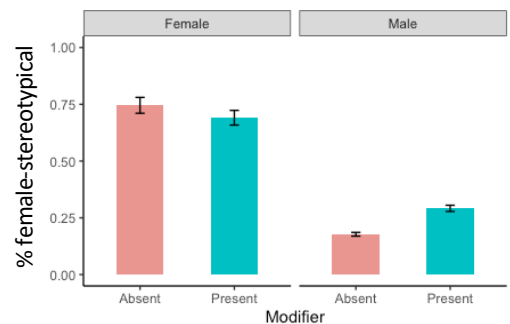


Figure 2. Female-stereotypical responses by Gender of the name (Female vs. Male) and Modifier (Absent vs. Present)

Analyzing simple effects, we found that (1) with male names, there was significant difference ($\beta = 0.648$, $SE = 0.266$, $z =$

2.434, $p < .05$) in the proportion of female-stereotypical responses between two levels of Modifier, but (2) with female names, there was no such difference ($\beta = -0.322$, $SE = 0.208$, $z = -1.548$, $p = 0.122$).

We also observed a numerically greater increase in the proportion of gender-astereotypical responses among male participants (Figure 3, right) than among female participants (Figure 3, left). However, we did not run any formal statistical tests since they might be unreliable given the difference in the number of participants between the two genders.




Figure 3. Female-stereotypical responses of female participants (*top*) and female-stereotypical responses of male participants (*bottom*)

Interim Discussion

Experiment 1 provided support for informativity expectations, demonstrating that the inclusion of a gender modifier shifts occupation-related predictions towards a somewhat gender-stereotypical direction.

However, an offline experiment like Experiment 1 only shows the end results of language processing under no time pressure. It is possible that the effects of informativity expectations emerge differently in an online experiment where participants are required to respond as fast as they can. In Experiment 2, we are interested in the real-time processing profile of informativity expectations. In addition, we have two other goals in the second experiment. Firstly, our goal is to ensure a balanced consideration of both stereotypical and non-stereotypical occupations. This is achieved by presenting participants with one of the two occupation options in each

trial and having it compared with an unrelated option. Second, we aim to examine informativity expectations with a reading-time based method, which complements the findings based on other online methods such as visual world eye-tracking.

Experiment 2: A-Maze Task

In Experiment 2, we test our predictions using real-time RTs. If occupational gender stereotypes affect sentence processing, we predict that gender-stereotypical occupations will yield faster RTs than gender-astereotypical occupations, at least when the gender modifier is absent. If informativity plays a role (but does not override the influence of gender stereotypes) in sentence processing, we predict a smaller difference in RTs between gender-stereotypical and gender-astereotypical occupation nouns when gender modifier is present than when it is absent.

Participants We recruited 120 participants (77 female, 43 male) online. All participants self-identified as native speakers of Mandarin Chinese and as currently residing in Mainland China.

Design & Materials Experiment 2 used the A-Maze method (Boyce, Futrell, & Levy, 2020), in which participants read one word at a time by choosing between the correct continuation and an incorrect one (Figure 4). When an incorrect continuation was chosen, participants would receive an immediate feedback message telling them to try again (with a penalty of 1000ms). They could only proceed in a trial when they clicked the correct continuation. Incorrect continuations were automatically selected by a neural network language model. Following Levison et al. (2023), we used multilingual BERT (Devlin et al., 2019) to generate incorrect continuations that had very low contextual probability but matched the correct ones in terms of frequency, and then we hand-corrected them.

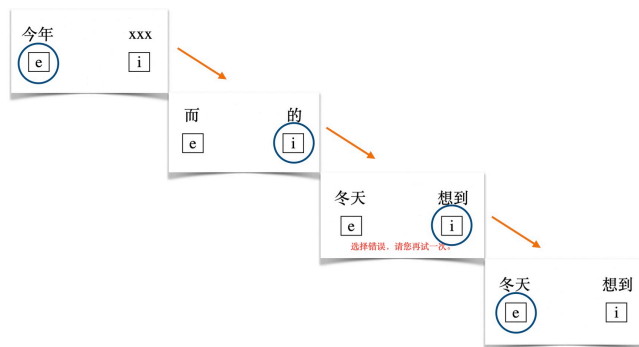


Figure 4. An illustration of the A-Maze task. Participants choose between two options, only one of which is the correct continuation for the sentence. If the correct choice is made, the following screen will automatically show up. When the wrong choice is made (as Screen 3 shows),

participants would receive a prompt that says, “Please try again”.

We manipulated Modifier (presence vs. absence) the way we did in Experiment 1. Instead of manipulating Gender of the proper name, we manipulated Stereotypicality by varying whether the occupation given in the correct continuation and the gender indicated by the proper name are stereotypically related. Modifier (present vs. absent) and Stereotypicality (stereotypical vs. astereotypical) were crossed in a within-participant 2x2 design (Table 2). Again, when the gender modifier was present, it was consistent with the gender of the name. The same pre-normed names and occupations from Experiment 1 were used here. As in Table 1, *Wang Hong* is an example of a typical female name and *Wang Wei* is an example of a typical male name.

Table 2. Example target items in Experiment 2

	STEREO-TYPICALITY	MODIFIER
a.	stereotypical Wang Hong shi yi-ge baomu. (“Wang Hong is a babysitter.”)	absent
b.	astereotypical Wang Wei shi yi-ge baomu. (“Wang Wei is a babysitter.”)	absent
c.	stereotypical Wang Hong shi yi-ge nv baomu. (“Wang Hong is a female babysitter.”)	present
d.	astereotypical Wang Wei shi yi-ge nan baomu. (“Wang Wei is a male babysitter.”)	present

We created 24 sets of target items and 36 fillers. We excluded the attention check questions from the 48 fillers we used in Experiment 1, because we can evaluate accuracy rates of correct continuations in the first attempts per word region per trial. We also modified the other fillers used in Experiment 1 to better suit the A-maze task. Following Experiment 1, four lists were created using a Latin square design to make sure that each participant saw each item in only one condition.

Procedure The experiment was hosted online through PCIBEX (Zehr & Schwarz, 2018; <https://www.pcibex.net/>). Each participant was randomly assigned one of the four lists. Item presentation was randomized. At the start of the experiment, participants were given instructions for the A-Maze task and completed practice trials. Participants were told that they should base their choices on whether they were good continuations for a natural sentence. On average, the experiment took about 12 minutes.

Analysis

Exclusion One participant failed to pass the 80% accuracy threshold and was removed from all analyses. We removed

RTs below 100ms or above 5000ms. We also excluded trials wherein participants made an incorrect choice before or at the critical word (i.e., the occupation noun). The exclusion criteria led to 10% loss of the target trials data.

Statistical Model We fitted linear mixed effects models predicting the log-transformed RTs of the occupation noun regions. We included Modifier and Stereotypicality as fixed effects, and random intercepts for participants and items. The inclusion of random slopes followed the same procedure described in Experiment 1. We report the maximal model that allowed for convergence. When a significant interaction was detected, we performed follow-up analyses. The analysis was carried out in R using *lme4* (Bates, Mächler, Bolker, & Walker, 2015). Significance was determined using Satterthwaite’s method of approximating degrees of freedom in the *lmerTest* package (Kuznetsova, Brockhoff, & Christensen, 2017). The significance level was set at .05.

Results

Visual inspection of the data (see Figure 5) shows three patterns that we predicted. First, when the gender modifier is absent, gender-astereotypical occupation nouns induced longer RTs than gender-stereotypical occupation nouns. Second, the presence of the gender modifier reduced the RTs at the occupation noun regions. Third, when the gender modifier is present, the difference in RTs between the two levels of Stereotypicality became much smaller. Our statistical analyses confirmed these observations. Results showed a main effect of Stereotypicality ($\beta = 0.043$, $SE = 0.013$, $t = 3.292$, $p = .003$) and a main effect of Modifier ($\beta = -.068$, $SE = .012$, $t = -5.5065$, $p < .001$). We also found a significant interaction between Stereotypicality and Modifier ($\beta = -.049$, $SE = .025$, $t = -1.966$, $p = .049$). Our follow-up nested comparisons suggest (1) when the gender modifier was absent, stereotypical occupations were processed faster than astereotypical occupations ($\beta = 0.067$, $SE = 0.018$, $t = 3.799$, $p < .001$), but (2) when gender modifier was present, levels of Stereotypicality were not different from each other ($\beta = 0.017$, $SE = 0.017$, $t = 0.986$, $p = .324$).

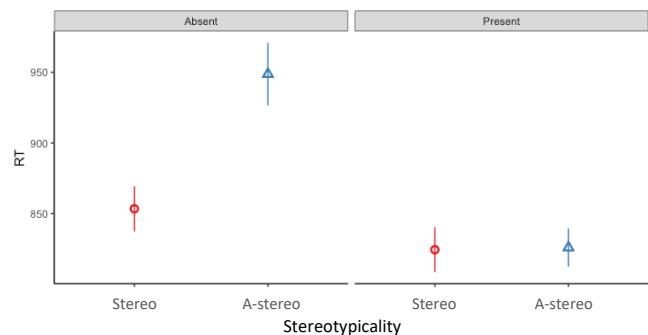


Figure 5. RTs (in ms) at the occupation noun region by Modifier (absent vs. present) and Stereotypicality (stereotypical vs. astereotypical)

Interim Discussion

Experiment 2 found similar evidence that informativity expectations can guide prediction, whereby the presence of the gender modifier facilitates the processing of gender-astereotypical occupation. This suggests that effects of informativity exist, but effects of gender stereotypes are hard to suppress, echoing our findings in Experiment 1.

General Discussion

The two studies presented here investigated informativity expectations. We showed that not only is world knowledge integrated in language processing but also expectations for newsworthiness affect language predictions (Rohde & Rubio-Fernandez, 2022; Rohde, Futrell, & Lucas, 2021; Brown & Dell, 1987; Grice, 1975, etc.). We used the optional gender modifier in Mandarin Chinese as a test case. When the gender of a person is known, overt use of a matching gender modifier should be regarded as redundant information unless, in our case, it introduces a gender-astereotypical occupation.

Results of both studies showed a clear impact of informativity expectations in guiding offline sentence completion and as well as real-time next-word prediction. When the gender modifier is present, we have observed either reduced preference for gender-stereotypical occupation responses or reduced RTs for gender-astereotypical occupation nouns. Our findings add support to the role of informativity expectations in facets of language processing.

However, it is also evident in our data that informativity expectations for non-redundant use of the gender modifier did not override the influence of gender stereotypes. This is somewhat unexpected given that prior work (e.g., Rodhe et al., 2022) has shown that informativity expectations steered comprehenders' preferences away from typical scenarios to less typical ones, e.g. from 'yellow banana' to 'yellow notebook'. Comparing our studies with theirs, one immediately noticeable difference emerges. What we tested was socially enforced stereotypes in gender roles and what they tested was prototypical physical properties in objects. The divergence between our findings might suggest a fundamental difference between stereotypes and prototypes from a novel angle. Stereotypes can be harder to resist and overcome, even when there are other constraints such as informativity expectations that have an opposing force. This has been hinted at in von der Malsburg et al.'s (2016) work on pronoun choice and gender stereotypes, where co-referring *president* with female pronoun *she* still incurred processing difficulty despite having a female presidential candidate in the country. Our observations are also reminiscent of Papineau et al.'s (2022) reading times data, wherein gender ideologies did not reliably predict the processing of gender-neutral roles. Linguistic studies like theirs and ours can be a useful tool for tapping into domains of human cognition, including but not limited to gender stereotypes.

Our empirical investigations have also revealed other interesting results that can inspire future work. The response

data from Experiment 1 presented an asymmetry between genders of names in that the effects of informativity were not pronounced in conditions where female names were used. One possible explanation is that the absence and presence of the female gender modifier are similarly frequent before a female-stereotypical occupation. To address the potential effects of frequency, we will have to run a corpus analysis to compare frequencies. In the meantime, while we have yet to test for the differences observed in results between female and male participants, individual factors such as gender appear to have influenced these variations. Notably, the influence of participants' gender was particularly evident in conditions involving the use of female names. With the presence of the gender modifier, male participants gave their occupation responses almost by chance, meaning that they were sensitive to the presence of the female gender modifier. However, the presence of the gender modifier did not trigger female participants to reason about the possibly redundant information. If this is related to frequency as well, it might suggest that female participants are much more frequently exposed to the combination of *gender modifier* + *female-stereotypical name*. If this is related to gender ideologies, however, our study could be an interesting first step into an individual differences study.

Conclusion

The present study showed that information about gender triggers occupational gender stereotypes. We found clear evidence that supports the role of informativity expectations in language processing; nevertheless, the impact of gender stereotypes was proved difficult to overcome, echoing findings in previous work on the interaction between language and gender stereotypes. Our work also extended investigation of informativity expectations into a socially situated context, a domain that had primarily been focused on referent communication and identification. This suggests that informativity expectations may be generally applicable in various types of contexts.

References

- Boyce, V., Futrell, R., & Levy, R. P. (2020). Maze Made Easy: Better and easier measurement of incremental processing difficulty. *Journal of Memory and Language*, *111*, 104082.
- Brown, P. M., & Dell, G. S. (1987). Adapting production to comprehension: The explicit mention of instruments. *Cognitive psychology*, *19*(4), 441-472.
- Engelhardt, P. E., Bailey, K. G., & Ferreira, F. (2006). Do speakers and listeners observe the Gricean Maxim of Quantity?. *Journal of memory and language*, *54*(4), 554-573.
- Grice, H. P. (1975). Logic and conversation. *Syntax and Semantics*, *3*, 41-58.
- McRae, K., & Matsuki, K. (2009). People use their knowledge of common events to understand language, and

- do so as quickly as possible. *Language and linguistics compass*, 3(6), 1417-1429.
- Papineau, B., Podesva, R., & Degen, J. (2022). 'Sally the Congressperson': The Role of Individual Ideology on the Processing and Production of English Gender-Neutral Role Nouns. *In Proceedings of the Annual Meeting of the Cognitive Science Society (Vol. 44, No. 44)*.
- Rohde, H., & Rubio-Fernandez, P. (2022). Color interpretation is guided by informativity expectations, not by world knowledge about colors. *Journal of Memory and Language*, 127, 104371.
- Rohde, H., Futrell, R., & Lucas, C. G. (2021). What's new? A comprehension bias in favor of informativity. *Cognition*, 209, 104491.
- Rubio-Fernandez, P. (2019). Overinformative speakers are cooperative: Revisiting the Gricean Maxim of Quantity. *Cognitive science*, 43(11), e12797.
- Rubio-Fernández, P. (2016). How redundant are redundant color adjectives? An efficiency-based analysis of color overspecification. *Frontiers in psychology*, 7, 153.
- Sedivy, J. C. (2003). Pragmatic versus form-based accounts of referential contrast: Evidence for effects of informativity expectations. *Journal of psycholinguistic research*, 32, 3-23.
- Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268(5217), 1632-1634.
- von der Malsburg, T., Poppels, T., & Levy, R. P. (2020). Implicit gender bias in linguistic descriptions for expected events: The cases of the 2016 United States and 2017 United Kingdom elections. *Psychological science*, 31(2), 115-128.
- Feigenbaum, E. A. (1963). The simulation of verbal learning behavior. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought*. New York: McGraw-Hill.
- Hill, J. A. C. (1983). A computational model of language acquisition in the two-year old. *Cognition and Brain Theory*, 6, 287-317.
- Matlock, T. (2001). *How real is fictive motion?* Doctoral dissertation, Psychology Department, University of California, Santa Cruz.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Ohlsson, S., & Langley, P. (1985). *Identifying solution paths in cognitive diagnosis* (Tech. Rep. CMU-RI-TR-85-2). Pittsburgh, PA: Carnegie Mellon University, The Robotics Institute.
- Shrager, J., & Langley, P. (Eds.) (1990). *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
- Zehr, J., Schwarz, F. (2018). PennController for Internet Based Experiments: IBEX. <https://doi:10.17605/OSF.IO/MD832>