Contrastive inferences are sensitive to informativity expectations, adjective semantics and visual salience

Paula Rubio-Fernandez^{1,2*}, Helena Aparicio², Vishakha Shukla³ & Julian Jara-Ettinger⁴

¹Department of Philosophy, University of Oslo; ²Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology; ³Dr Shroff's Eye Hospital in Delhi; ⁴Department of Psychology, Yale University. *Correspondence to: prubio@mit.edu

Word count: 5105

Abstract

People sometimes derive contrastive inferences from adjective-modified noun phrases. For example, the description ‘the short pencil’ would normally be understood to contrast a shorter and a longer pencil. In a series of eye-tracking studies, Sedivy (2003, 2004) found that scalar and material adjectives elicited contrastive inferences, but color adjectives did not, which she interpreted as a difference in informativity expectations. Here we assessed whether not only pragmatics, but also semantic and perceptual factors contribute to the derivation of contrastive inferences. We used Sedivy’s eye-tracking paradigm to compare the interpretation of scalar, material and color adjectives, employing new analyses to test the view that different fixation patterns should be observed on the target and the competitor objects depending on the semantics of the adjective and the visual salience of the encoded property. The results support our hypotheses, confirming that all three adjective types can elicit contrastive inferences, but their derivation is shaped differently by pragmatic, semantic and perceptual factors.

Keywords: scalar and non-scalar adjectives; contrastive and descriptive uses; color words; contrastive inferences; pragmatics; semantics; visual salience.

Introduction

Adjectives can be used contrastively to distinguish an intended referent from other objects of the same kind. For example, the description ‘the short pencil’ would normally be used to contrast a short pencil with a longer one. Sedivy and colleagues (1999) used the visual-world paradigm to tap into the derivation of this type of meaning or *contrastive inference*. In their study, participants were asked to move objects in a display following descriptions with scalar adjectives, such as ‘the short pencil.’ Critically, some trials had a *contrast object* (e.g., a longer pencil), supporting a contrastive interpretation of the adjective, whereas other trials did not, only allowing for a descriptive interpretation (compare displays (1a-1b) in Fig. 1). Contrastive interpretation was measured by comparing fixations on the *target object* relative to fixations on a *competitor object* (e.g., short scissors), which did not have a contrast in the display. Thus, when processing ‘Move the short...’ in display (1a), participants who derive a contrastive inference should anticipate the target before hearing the noun. That is, in hearing ‘short’, participants should fixate on the short pencil (because there is a long pencil in the display) and not on the short scissors (because there are no long scissors in the display). However, in the no-contrast condition in (1b), the temporary ambiguity created by ‘short’ could only be resolved when hearing the noun, with participants fixating comparably on the target and the competitor objects while processing the adjective.

Sedivy et al. (1999) observed that participants derived contrastive inferences when interpreting scalar adjectives, identifying the target earlier when there was a contrast object in the display than when the adjective was used descriptively. Sedivy (2004) used the same eye-tracking paradigm and found contrastive effects with scalar and material adjectives, but not with

color adjectives. She interpreted this difference as an effect of the varying frequency with which these adjectives were used descriptively in a language production task: participants produced color adjectives in non-contrastive contexts around half the time, but rarely did so with material and scalar adjectives. That is, participants often mentioned the color of single objects of a kind, but not other properties such as their material, size or length. Sedivy (2003, 2004) argued that the derivation of contrastive inferences is driven by general expectations of informativity (Grice, 1975): when a speaker produces a material or scalar adjective, the listener is entitled to assume that the adjective has a contrastive function because that is how those adjectives are most commonly used. However, this expectation is not warranted when a speaker produces a color adjective, because those are also used in non-contrastive contexts.

Fig. 1: Sample displays from the three adjective types in the two experimental conditions. The description of the target (which was presented auditorily) is written under each display.

Supporting the informativity account of contrastive inferences, Sedivy (2003) showed that, when describing objects with predictable colors (e.g., a yellow banana), participants used color adjectives only when there was a contrast object in the display (e.g., a brown banana). Moreover, participants were faster to identify 'the yellow banana' when a brown banana made the color adjective contrastive, relative to when the predictable color was purely descriptive. Sedivy (2003) concluded that referential contrast effects are triggered by listeners' expectations of informativity, therefore being pragmatic in nature. In her view, the *default description* of objects with unpredictable colors (e.g., cups) includes color adjectives and therefore, their use does not elicit contrastive inferences. However, the default description of objects with predictable colors (e.g., bananas) omits color adjectives, and their use does elicit contrastive inferences. Like predictable colors, scalar and material adjectives are not part of the default description of objects and are therefore interpreted contrastively.

Sedivy's (2003, 2004) results support the informativity view of contrastive inferences, but this account faces some challenges. First, speakers used descriptive color adjectives less than half of the time when naming objects with unpredictable colors (Sedivy, 2004), making it unlikely that listeners expect color words to be used by default. Moreover, Rubio-Fernandez (2019) recently observed that speakers' tendency to use color adjectives descriptively varies greatly depending on the discriminability of the referent, the density of the display and the visual salience of color. Rubio-Fernandez (2016) also observed that the lexical category of the noun affected color descriptions, with clothes eliciting descriptive color adjectives around 90% of the time whereas geometrical shapes did so around 40% of the time. Finally, Rubio-Fernandez

(2019) observed that size adjectives are used descriptively more than 60% of the time when their absolute size pops out in a display. Rubio-Fernandez (2016, 2019) interpreted this variability as a result of the varying efficiency of descriptive modification in different contexts. Overall, the high variability in the production of descriptive color adjectives challenges the view that these are part of how we refer to objects by default.

Moreover, Aparicio et al. (2016; Aparicio, 2018) recently compared the interpretation of color and scalar adjectives using Sedivy's eye-tracking paradigm with geometrical shapes, for which colors are not predictable, and found referential contrast effects with both types of adjectives, contrary to what Sedivy (2004) had observed with everyday artefacts. However, rather than using the adjective and noun as time windows for analysis (as Sedivy had done), Aparicio et al. split the phrase into three short windows of 150ms each. Methodological differences such as the type of visual stimuli or the time windows used for analysis could therefore explain the different results observed with color adjectives in these two studies.

Aparicio et al. (2016; Aparicio, 2018) observed that participants disambiguated the target from the competitor in the same time window during the processing of color and scalar adjectives, but were faster to attribute color than scalar properties (i.e. participants identified the two color-matching objects faster than the two scalar-matching objects). Aparicio and colleagues attributed these different looking patterns to the different semantics of these adjectives. While scalar adjectives are interpreted in relation to a comparison class, color adjectives are not, because they encode a non-gradable property (e.g., a pencil may be short compared to a longer pencil in the context, or in relation to the average length of pencils, but it can be yellow in and of itself; see Kennedy, 1999, 2007; Kennedy & McNally, 2005). Therefore,

scalar adjectives are more context dependent than color adjectives, possibly incurring higher processing costs.

Another important difference between color and scalar adjectives is the relative visual salience of the properties they encode, with color being potentially more salient than scalar properties. Hence, not only semantic but also perceptual factors could explain the different looking patterns observed by Aparicio et al. (2016; Aparicio, 2018) between scalar and color adjectives. The role of semantic and perceptual factors in the derivation of contrastive inferences could be teased apart by comparing the interpretation of color, material and scalar adjectives, since color and material are non-scalar, but material is not as visually salient as color (Jara-Ettinger & Rubio-Fernandez, in preparation). Here we used Sedivy's eye-tracking paradigm with these three adjective types to investigate whether all contrastive effects are pragmatic in nature (as Sedivy argued), or whether the semantics and perceptual properties of different adjectives also affect their interpretation.

Contrastive effects in Sedivy's eye-tracking paradigm

In her eye-tracking studies, Sedivy (2003, 2004) used a *target advantage measure*, calculated as the proportion of time spent looking at the target (e.g., the short pencil in display (1a) in Fig. 1) minus the proportion of time spent looking at the competitor (e.g., the short scissors in the same display) during three consecutive time windows corresponding with the adjective, the noun, and 10 video frames after the noun. If the adjective was interpreted contrastively, the target advantage was expected to be greater in the contrast than the no-contrast condition. Similarly, Aparicio (2018), compared fixations on the target and the competitor object in three consecutive

windows of 150ms in order to time the disambiguation of the instruction (i.e. when in processing ‘the short pencil’ participants started to show a preference for the short pencil over the short scissors). In addition, Aparicio (2018) directly compared fixations on the target object across the two experimental conditions in order to further investigate whether an earlier disambiguation in the contrast condition corresponded with an increase in target fixations (rather than a reduction in competitor fixations).

Here we further explored the role of the competitor object in the derivation of contrastive effects. In particular, we analyzed fixations on the target and the competitor separately in order to tease apart the two possible effects that a contrastive inference might have in this paradigm: namely, anticipating the target (what we will call a *target boost*) or reducing the interference of the competitor object (what we will call a *competitor reduction*). Based on previous findings, we expect that all adjective types will elicit contrastive inferences. However, we predict different processing blueprints for color, material and scalar adjectives based on their semantics and the visual salience of the properties they encode.

Given their relational semantics, we predict that scalar adjectives will reveal a strong target boost when interpreted contrastively, provided the interpreter conceptualizes the two objects of a kind as contrasting along the relevant dimension. For example, during the preview of display (1a) in Figure 1, participants should conceptualize the two pencils as a short and a long pencil in order to anticipate the noun ‘pencil’ when processing the adjective ‘short.’ Relatedly, we do not predict a strong competitor reduction because attributing a scalar property requires a comparison class. Thus, since the scissors do not have a contrast object in the display, they should not be conceptualized as short during the preview window. Therefore, when processing

scalar adjectives, competitor objects should exert comparable interference in the contrast and no-contrast conditions.

Unlike scalars, color and material adjectives encode non-gradable properties that can be attributed without a comparison class. That means that the competitor object may interfere during processing: when participants hear ‘the black...’ in display (2a) or ‘the leather...’ in (3a), they may consider the black chair and the leather gloves, even if they are single objects. We therefore predict that if color and material adjectives are interpreted contrastively, the pragmatic inference should reduce the competitor interference in the contrast condition relative to the no-contrast condition. That is, participants should look less at the black chair than the black suitcase in displays (2a-2b), and less at the leather gloves than the leather boots in displays (3a-3b).

Even if interpreting color adjectives does not require a comparison class, we predict that when color adjectives are interpreted contrastively, they will elicit a target boost because color contrast is perceptually salient. Therefore, during the preview of display (2a), participants should readily perceive the color contrast between the two lamps, which would prepare them to anticipate the noun when hearing ‘the black...’. Material adjectives, on the other hand, do not encode equally salient visual properties (Jara-Ettinger & Rubio-Fernandez, in preparation). Therefore, during the preview window, the two pairs of shoes in display (3a) may not be encoded as contrasting in material, preventing participants from anticipating the noun ‘shoes’ when hearing ‘the leather...’. We therefore hypothesize that the contrastive interpretation of

color adjectives, but not of material adjectives, should result in a target boost, even though both types of adjective are non-gradable.¹

Two main predictions follow from the above analysis: first, when interpreted contrastively, both scalar and color adjectives should reveal a target boost, the former because of their relational semantics and the latter because of the visual salience of color contrast. Material adjectives, on the other hand, should not reveal a significant target boost because material contrast is not as visually salient. Second, because color and material adjectives encode non-scalar properties, they should reveal a competitor reduction when interpreted contrastively. In the case of scalar adjectives, the competitor object may not be attributed the relevant property in the absence of a contrast, resulting in comparable interference in the two conditions. More generally, we argue that Sedivy's eye-tracking paradigm has the potential to reveal the nature of contrastive inferences (i.e. whether they are affected by pragmatic, semantic and/or perceptual factors), if fixations on the target and competitor objects are analyzed separately.

Experiment

Key to our predictions is the assumption that in Sedivy's eye-tracking paradigm, participants need to conceptualize the two objects of the same kind as contrasting along a certain dimension (e.g., two pencils of different length or two lamps of different colors, as in Fig.1) in order to be able to anticipate the noun when processing the adjective in the critical condition. Therefore, in

¹ We consider the preview window to be crucial for the derivation of contrastive inferences in this paradigm, but we do not make any predictions about participants' fixation patterns during that window. That is, we do not think that participants would look differently at the two objects depending on whether or not they realize that they are different colors, or made of different materials. However, while this conceptualization process may not reveal different looking behavior during the preview window, it should allow participants to show anticipatory looking during the processing of the instruction – which is what we are predicting.

our study, the preview window was relatively long (3s). The instructions, on the other hand, included only a description of the target (e.g., ‘the short pencil’) and were therefore shorter than in previous studies (e.g., ‘Pick up the short pencil’ in Sedivy, 2004), potentially resulting in a delayed anticipatory effect. Given the short instructions, and following a previous study using the same paradigm and visual materials in four different languages (Rubio-Fernandez & Jara-Ettinger, under review), we focused our statistical analyses on the noun window.

Our analyses are novel, comparing fixations on the target and on the competitor separately across the two conditions for each adjective type (rather than calculating a differential score between the two objects). The point of doing separate analyses for the target and the contrast object was to test whether the contrastive interpretation of the adjective in the critical condition resulted in a target boost and/or a competitor reduction depending on the adjective type.

Finally, the two experimental conditions were split into two separate trial blocks, such that the contrast condition was administered in the first trial block, and the no-contrast condition in the second. This design was intended to maximize pragmatic reasoning as it has been observed that mixing contrastive and descriptive adjectives prevents participants from deriving contrastive inferences (Grodner & Sedivy, 2011). In other words, if the speaker in the task used adjectives not only to preempt an ambiguity between two objects of the same kind, but also to describe single objects, then participants would not have pragmatic grounds to anticipate the noun in contrastive trials. For this reason, the contrast condition was always presented in the first block of trials, before the no-contrast condition.

Methods

Participants

Sixty participants were recruited for the study from the participant pool at the Department of Brain and Cognitive Sciences at MIT. Participants included MIT students and students from other universities in the Cambridge area, as well as local residents. All participants were native speakers of English and received monetary compensation for their time.

Materials and procedure

A list of 8 color adjectives (black, blue, brown, green, orange, red, white and yellow), 8 material adjectives (cotton, glass, gold, leather, metal, paper, plastic, wooden and woolen) and 8 scalar adjectives (large, narrow, short, small, tall, thick, thin and wide) were compiled for the study.

Two displays of four objects were created for each adjective, one for each experimental condition. In the No-Contrast condition, the display included an object that shared the relevant property with the target, whereas in the Contrast condition, the display included both an object that shared the relevant property with the target, and an object of the same kind as the target. For sample items, see Figure 1. A set of 4 displays were used as warm-up trials and another set of 48 displays (24 with competitors and 24 without competitors) were used as fillers. Warm-up and filler trials were not analyzed.

Target descriptions were recorded by a native speaker of English who did not stress the adjectives contrastively. Unlike in the studies by Sedivy (2003, 2004) and Aparicio et al. (2016; Aparicio, 2018), target descriptions were not embedded in a full instruction (i.e. participants heard ‘the short pencil’ rather than ‘pick up the short pencil’ or ‘click on the short pencil’).

Because the instructions were overall shorter than in previous studies, we expected the contrastive effects to be delayed.

The materials were presented in two blocks of trials, with a 10s break in between the two blocks. In the first block, all displays contained two objects of the same kind and were comprised of the 24 trials from the Contrast condition plus 24 fillers. The target was one of the two objects of the same kind in the experimental condition and one of the other two objects in the filler trials. The target description included an adjective in the Contrast condition, but not in the filler trials (e.g., ‘the short pencil’ was used in a display with two pencils and ‘the elephant’ in a display with a single elephant). The filler trials were designed to prevent participants from anticipating the target in critical trials. In the second block of trials, all four items in each display were different types of objects, but two of the objects shared a property. The target was always described using a descriptive adjective. As in the first block, the filler trials were designed to prevent participants from anticipating the target in critical trials. The 24 adjectives selected for the study were used once in the first block of trials and twice in the second; however, the 24 adjectives modified 72 different nouns in order to prevent participants from anticipating the target when the adjectives were repeated in the second block of trials. Trial presentation was randomized individually within each block.

Participants had 3s to preview the display before the description of the target started. The preview was relatively long to allow participants to compare the objects and conceptualize the relevant contrast. From the onset of the target description, participants had 3s to click on the target. In between trials, participants had to click on a cross in the center of the screen to ensure that the cursor was equidistant from the four objects at the start of each trial.

Participants were asked to inspect the display during the preview window and click on the target object as fast and accurately as they could on hearing its description.

Eye movements were recorded with a portable eye-tracking system (RED-m by SMI) that measured eye position at a sampling rate of 120 Hz and had a spatial resolution (RMS) of 0.1° and an accuracy of 0.5°. Participants were seated about 60 cm from the computer screen, with a contact-free set up. The task lasted approximately 20min.

Results

Percentage of fixations on the four objects in a display are plotted in Figures 2-4 by adjective type and condition. Proportion of time spent fixating on the target and the competitor were calculated separately in the Contrast and No-Contrast conditions between the noun onset and offset (see Fig. 5). Noun onset and offset were standardly corrected by +200ms to adjust for the time it takes to launch a saccade. The proportion of target fixations were expected to be higher in the Contrast condition than in the No-Contrast condition for scalar and color adjectives, revealing a target boost. The proportion of competitor fixations were expected to be lower in the Contrast condition than in the No-Contrast condition for color and material adjectives, revealing a competitor reduction.

Fig. 2: Percentage of time spent fixating on the four objects in a display during the processing of the target description. Vertical lines mark the average noun onset and offset.

Fig. 3: Percentage of time spent fixating on the four objects in a display during the processing of the target description. Vertical lines mark the average noun onset and offset.

Fig. 4: Percentage of time spent fixating on the four objects in a display during the processing of the target description. Vertical lines mark the average noun onset and offset.

Fig.5: Proportion of time spent fixating on the target (left) and on the competitor (right) during the noun window for each adjective type in the two experimental conditions. Error bars represent 95% bootstrapped confidence intervals.

Data were pre-processed in order to carry out two regression analyses, one on the target fixations and another on the competitor fixations. The 24 adjectives used in the study were repeated in the two conditions, but the nouns varied across conditions, with the noun window

ranging from 289ms/1 syllable to 828ms/4 syllables ($M = 553\text{ms}/1.4$ syllables). The length of the noun windows was not significantly different in the Contrast condition ($M=557\text{ms}$) and the No-Contrast condition ($M=549\text{ms}$; $p = .832$), but the overall variability introduced high item variability in the two eye-tracking measures. To reduce this variability, proportions of target fixations and contrast fixations were averaged for each participant across the 8 nouns used for each adjective type in each condition. Because noun duration is a theoretically-motivated time window, this solution was preferred over selecting a time window of arbitrary length that was comparable across items. It must also be noted that contrastive effects were not expected in the adjective window (cf. Sedivy, 2003, 2004; Aparicio et al., 2016) because the instructions were shorter than in previous studies (see also Rubio-Fernandez & Jara-Ettinger, under review).

We used linear mixed effects regression to model each eye-tracking measure (average proportion of time fixating on the target and on the competitor) as the outcome variable with Condition (Contrast/No Contrast) and Adjective (Color/Material/Scalar) as fixed effects. Participants were included as random effects, and random slopes were included for Condition and Adjective (but without their interaction because that model did not converge). Given our experimental hypotheses, we used a different reference level for each model, corresponding to the adjective that was predicted to behave differently from the other two: for the Target model, the Material/No-Contrast condition was dummy coded as the reference level, whereas for the Competitor model, it was the Scalar/No-Contrast condition. Each model was fit using the *lme4* package (Bates et al., 2014) under the statistical computing language *R* (Version 1.2.1335; R Core Team, 2018). Model outputs are reported in Table 1.

In the Target model, the Condition x Adjective interaction was significantly different in the Material and Scalar conditions ($\beta = 0.0660$; $t = 2.553$) and the Material and Color conditions ($\beta = 0.0842$; $t = 3.255$)². Comparing models with and without the interaction term between Condition and Adjective (holding random-effects structure constant) revealed a significant difference in favour of the model with an interaction term ($\chi^2(2) = 11.556$, $p = .0030$).

In the Competitor model, the Condition x Adjective interaction was significantly different in the Scalar and Color conditions ($\beta = -0.0605$; $t = -2.457$), but did not reach significance when comparing the Scalar and Material conditions ($\beta = -0.0392$; $t = -1.591$). Comparing models with and without the interaction term between Condition and Adjective (holding random-effects structure constant) revealed a significant interaction in favour of the model with an interaction term ($\chi^2(2) = 6.2083$, $p = .0448$).

Measure	Conditions	Beta	Standard Error	t value
Proportion Target fixations	Intercept	0.3647	0.0177	20.654
	Color	-0.0375	0.0193	-1.947
	Scalar	-0.0775	0.0185	-4.188
	Contrast	0.0497	0.0202	2.457
	Color:Contrast	0.0842	0.0259	3.255
	Scalar:Contrast	0.0660	0.0259	2.553
Proportion Competitor fixations	Intercept	0.2708	0.0146	18.494
	Color	0.0898	0.0181	4.958
	Material	0.0362	0.0176	2.052
	Contrast	-0.0285	0.0199	-1.435
	Color:Contrast	-0.0605	0.0246	-2.457
	Material:Contrast	-0.0392	0.0246	-1.591

Table 1: Model outputs for each eye-tracking measure.

² When interpreting *lmer* model outputs, we define significance as $|t| > 2$ (Gelman & Hill, 2007).

General discussion

Analysing participants' fixations on the target and the competitor separately revealed nuanced differences in the interpretation of the three adjective types. When interpreted contrastively, scalars elicited a target boost but no competitor reduction, on account of their relational semantics (Kennedy, 1999, 2001; Kennedy & McNally, 2005). Reversely, contrastive effects with material adjectives seemed to affect the competitor but not the target, which we predicted on the grounds that material is a non-scalar property. Color adjectives, on the other hand, elicited both a target boost and a competitor reduction, with the former being related to the visual salience of color contrast and the latter to their non-scalar semantics.

These results leave open the question of why in the original study, Sedivy (2004) did not observe contrastive effects with color adjectives (cf. Aparicio et al., 2016). Sedivy (2003, 2004) argued that contrastive inferences were triggered by expectations of informativity, with color adjectives failing to elicit contrastive inferences because they are often used descriptively (see Pechmann, 1989; Arts et al., 2011; Koolen et al., 2013; Rubio-Fernandez, 2016, 2019). In our study, the first block of trials only included contrastive adjectives, potentially increasing participants' expectations that color would be used canonically to distinguish between objects of the same kind (Grodner & Sedivy, 2011). This interpretation of our findings is consistent with Sedivy's informativity account.

However, it must be noted that the use of a two-block design introduced a limitation in the item analyses. Separating the two conditions in two trial blocks required that the adjectives accompanied two different nouns in each condition in order to avoid that participants could

anticipate the target in the second trial block. Using different nouns, however, introduced high item variability, which we tried to overcome by averaging the looking data by adjective and condition, resulting in weaker statistical analyses. Future studies that want to maximize pragmatic reasoning by using a two-block design should try to match the nouns in length and frequency across the two trial blocks in order to reduce item variability.

Bearing that limitation in mind, we conclude that contrastive inferences reveal different processing blueprints depending on (i) participants' expectations of informativity, (ii) the semantics of the adjective, and (iii) the visual salience of the encoded property. Therefore, rather than being a purely pragmatic phenomenon, contrastive inferences seem to be sensitive to multiple cognitive factors, including the semantics of the adjectives and their perceptual characteristics.

Above and beyond these factors, we assume that deriving a contrastive inference in this paradigm requires conceptualizing the relevant contrasts ahead of processing the description. That is, during the preview window, participants need to appreciate that there are two objects of the same kind that differ along a certain dimension. This requirement may cast doubts on the ecological validity of this paradigm. However, we must bear in mind that conceptualizing the relevant contrasts ahead of time is only required to *measure* the contrastive effect: in everyday conversation, we may appreciate that a speaker is using an adjective contrastively (or descriptively) at any point during interpretation, with the contrastive inference not always allowing us to anticipate the noun. What the results of this eye-tracking study show is that separate analyses of target and competitor fixations in Sedivy's paradigm can reveal how pragmatic, semantic and perceptual factors affect the derivation of contrastive inferences.

References

- Aparicio, H. (2018), *Processing context-sensitive expressions: The case of gradable adjectives and numerals* (Ph.D. thesis, University of Chicago).
- Aparicio, H., Xiang, M., & Kennedy, C. (2016), Processing gradable adjectives in context: A visual world study. *Semantics and Linguistic Theory*, 25, 413-432.
- Arts, A., Maes, A., Noordman, L., & Jansen, C. (2011). Overspecification in written instruction. *Linguistics*, 49, 555-574.
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and hierarchical/multilevel models*. New York, NY: Cambridge.
- Grice, H. P. (1975). Logic and conversation. In P. Cole and J. Morgan (Eds.), *Speech acts*. Academic Press (pp. 41-58).
- Grodner, D., & Sedivy, J. (2011). The effect of speaker-specific information on pragmatic inferences. In N. Pearlmuter and E. Gibson (Eds.), *The Processing and Acquisition of Reference*. MIT Press: Cambridge, MA (pp. 239-272).
- Jara-Ettinger, J., & Rubio-Fernandez, P. (in preparation). Modelling the effect of visual search on language production.
- Kennedy, C. (1999). *Projecting the adjective: The syntax and semantics of gradability and comparison*. New York: Garland. (1997 Ph.D. thesis, UCSD).
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30, 1-45.

- Kennedy, C., & McNally, L. (2005). Scale structure and the semantic typology of gradable predicates. *Language*, 81, 345–381.
- Koolen, R., Goudbeek, M., & Krahmer, E. (2013). The effect of scene variation on the redundant use of color in definite reference. *Cognitive Science*, 37, 395-411.
- Pechmann, T. (1989). Incremental speech production and referential overspecification. *Linguistics*, 27, 89-110.
- Rubio-Fernandez, P. (under review). Overinformative speakers are cooperative: Revisiting the Gricean Maxim of Quantity.
- Rubio-Fernandez, P. (2016). How redundant are redundant color adjectives? An efficiency-based analysis of color overspecification. *Frontiers in Psychology*, 7. Special Issue *Models of Reference*.
- Rubio-Fernandez, P., & Jara-Ettinger, J. (under review). Word order determines visual search and pragmatic reasoning: A cross-linguistic eye-tracking study.
- Sedivy, J. (2003). Pragmatic versus form-based accounts of referential contrast: Evidence for effects of informativity expectations. *Journal of Psycholinguistic Research*, 32, 3-23.
- Sedivy, J. (2004). Evaluating explanations for referential context effects: Evidence for Gricean mechanisms in online language interpretation. In J. Trueswell and M. Tanenhaus (Eds.). *Approaches to studying world-situated language use*. Cambridge, Massachusetts: MIT Press (pp. 345-364).
- Sedivy, J. C., Tanenhaus, M. K., Chambers, C. G., & Carlson, G. N. (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition*, 71, 109-147.

Acknowledgments

The authors thank Madeleine Long and Francis Mollica for helpful discussions of this work. This research was supported by a *Young Research Talent Grant* (Research Council of Norway, 230718) awarded to PRF and a *Google Faculty Research Award* to JJE.