Cognitive Science 43 (2019) e12797

© 2019 The Authors. Cognitive Science published by Wiley Periodicals, Inc. on behalf of Cognitive Science Society (CSS). All rights reserved.

ISSN: 1551-6709 online

DOI: 10.1111/cogs.12797

Overinformative Speakers Are Cooperative: Revisiting the Gricean Maxim of Quantity

Paula Rubio-Fernandez^{a,b}

^a*Department of Philosophy, University of Oslo*

^b*Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology*

Received 21 March 2019; received in revised form 1 October 2019; accepted 1 October 2019

Abstract

A pragmatic account of referential communication is developed which presents an alternative to traditional Gricean accounts by focusing on cooperativeness and efficiency, rather than informativity. The results of four language-production experiments support the view that speakers can be cooperative when producing redundant adjectives, doing so more often when color modification could facilitate the listener's search for the referent in the visual display (Experiment 1a). By contrast, when the listener knew which shape was the target, speakers did not produce redundant color adjectives (Experiment 1b). English speakers used redundant color adjectives more often than Spanish speakers, suggesting that speakers are sensitive to the differential efficiency of prenominal and postnominal modification (Experiment 2). Speakers were also cooperative when using redundant size adjectives (Experiment 3). Overall, these results show how discriminability affects a speaker's choice of referential expression above and beyond considerations of informativity, supporting the view that redundant speakers can be cooperative.

Keywords: Referential communication; Audience design; Redundancy; Informativity; Cooperativeness; Efficiency

1. Introduction

Redundancy provides a puzzle for pragmatic theories of referential communication. If there was a single star on the computer screen, why would I ask you to click on “the blue

Correspondence should be sent to Paula Rubio-Fernandez, Department of Philosophy, University of Oslo, Georg Morgenstiernes hus, Blindernveien 31, 0313 Oslo, Norway. E-mail: prubio@mit.edu

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

star”? Calling for a solution to this puzzle, experimental research has shown that people often use adjectives redundantly when referring to entities in their physical environment (e.g., Arts et al., 2011a; Engelhardt et al., 2006; Engelhardt & Ferreira, 2014; Maes et al., 2004; Sedivy, 2003, 2004; van der Sluis & Krahmer, 2007). Some authors suggest that redundancy is easier for the speaker who does not need to determine whether an adjective is necessary in the context (e.g., Belke, 2006; Belke & Meyer, 2002; Koolen et al., 2013; Pechmann, 1989), but it has also been argued that redundant descriptions may help the listener identify the intended referent (e.g., Arts et al., 2011b; Mangold & Pobel, 1988; Paraboni & van Deemter, 2014; Paraboni et al., 2007; Rubio-Fernandez, 2016; Sonnenschein & Whitehurst, 1982).

According to Grice’s (1975) Maxim of Quantity, speakers should provide their interlocutors with as much information as necessary, but not more. Standard pragmatic accounts in the Gricean tradition are therefore based on *informativity*: the amount of information that is required to establish unique reference in a given context. These accounts draw a three-way distinction between *underinformative*, *minimal*, and *overinformative referential expressions*. Key to this distinction are a referent’s *competitors*: objects of the same category found in the situational context (e.g., how many stars the listener must choose from when hearing “Click on the blue star”). Thus, the expression “the blue star” would be minimal if there were two stars of different colors on the screen; but the same expression would be underinformative if there were two blue stars, or overinformative if there were a single star.

In line with the Gricean view, the general assumption in the pragmatics literature is that minimal referential expressions that determine unique reference are preferred over underinformative or overinformative expressions (e.g., Davies & Katsos, 2010; Engelhardt et al., 2006, 2011; Engelhardt & Ferreira, 2014; Grodner & Sedivy, 2011; Heller et al., 2012; Pogue et al., 2016; Sedivy, 2003, 2004, 2007). I have previously challenged the traditional view by arguing that an optimal referential expression must take into account not only the referent and its competitors, but also the perceptual properties of the situational context (Rubio-Fernandez, 2016). Thus, in evaluating the request “Click on the blue star,” a pragmatic theory should consider not only how many stars are on the screen (which would determine the *informational value* of the description), but also how many blue shapes there are (which would determine the *discriminatory value* of the color adjective). While computational studies of referential expression generation have long considered the role of perceptual factors on the use of modification (e.g., Arts et al., 2011b; Koolen et al., 2013; Paraboni et al., 2007; Westerbeek et al., 2015), pragmatic accounts in the Gricean tradition have only been concerned with information quantity, measured in number of words.

In this paper, I will expand upon the account of redundant modification that I first proposed in Rubio-Fernandez (2016), according to which redundant modification can be efficient if it helps the listener identify the referent in the visual context (for an analogous proposal formalized in terms of entropy reduction, see Tourtouri et al., 2019). Here I will further argue that redundancy can be a cooperative choice if speakers are sensitive to the listener’s needs, and report four reference production experiments testing this view. Redundancy, however, has sometimes been confounded with lack of perspective taking, and efficiency with brevity. In the next two subsections, I will challenge the views that

redundancy is tantamount to egocentricity and efficiency to brevity, before turning to the question of what might be cooperative and efficient in referential communication.

2. Are redundant speakers egocentric?

Some researchers see the production of redundant modification as a speaker-internal process, whereas others see it as a listener-oriented process (Arnold, 2008). In other words, redundant modification can be understood as a form of *egocentric behavior* (benefitting the speaker) or a form of *audience design* (tailored to the listener). This debate has often been informed by *referential communication tasks*, in which a speaker has to request a target object from an array, ensuring they provide enough information for the listener to identify the intended referent. However, this kind of communicative situation in which speaker and listener can rely on the *co-presence* of their physical environment (Clark & Marshall, 1981) is not appropriate for the investigation of perspective taking, which requires that the interlocutors' perspectives be different (for discussion, see Barr, 2014; Brennan et al., 2010; Keysar, 1997). Because speakers and listeners normally rely on the same perceptual mechanisms, in situations of co-presence interlocutors share an assumption of *mutual salience*. That is, the speaker is entitled to assume that anything that is perceptually salient to them will also be salient for the listener, without having to adopt the listener's perspective. Since the debate on whether or not redundant speakers are egocentric has been informed by a paradigm that does not really require perspective taking, no study has been able to settle the matter after more than a decade of experimental work and theoretical discussion (for recent attempts, see Engelhardt & Ferreira, 2016; Tourtouri et al., 2019).

While inadequate for the investigation of perspective taking, situations of co-presence are a good test case for the relative efficiency of speaker-listener coordination (Clark & Wilkes-Gibbs, 1986). However, such an investigation requires drawing a distinction between *egocentric* and *efficient* processes—one that has often been blurred in the literature (e.g., Barr, 2014; Engelhardt & Ferreira, 2016; Tourtouri et al., 2019). Thus, in a situation where speaker and listener share a physical environment, the speaker can rely on their own visual perspective as part of their common ground with the listener. This, however, is not a form of egocentric behavior whereby the speaker is insensitive to the listener's perspective. Instead, relying on mutual salience is a relatively effortless (and hence efficient) way to be tuned to the interlocutors' shared perspective—what I shall call *visual common ground*.

Consider a situation where Harry and Sally are looking at the night skyline of New York City. If Harry said “Look at *that!*”, Sally would not have much to go by since “that” could effectively mean anything (see Clark et al., 1983). However, given their visual common ground, Harry was entitled to assume that the Empire State Building turning bright pink would be a remarkable sight not only for him, but also for Sally. To drive the point home, note how different this exchange would be if Harry had made the same remark while talking to Sally on the phone. Only in the latter situation could it be argued that Harry was egocentric: Situations of co-presence simply allow for more efficient speaker–listener coordination.

3. Can redundancy be efficient? A collaborative perspective

The present account assumes that reference is a collaborative process between interlocutors, as proposed by Clark and colleagues in their pioneering studies (Brennan & Clark, 1996; Clark, 1996; Clark & Marshall, 1981; Clark et al., 1983; Clark & Wilkes-Gibbs, 1986; Horton & Brennan, 2016; see also Ford & Olson, 1975; Olson, 1970). According to Clark's theory of language use, interlocutors collaborate in the accumulation of common ground during conversation. In this view, the minimal unit of joint action is the *contribution*, which consists of a presentation and an acceptance phase: While an utterance may count as a presentation, it needs to be accepted by the listener before the contribution can be considered complete. For example, a question is not accepted until the listener has produced a response that shows that they have understood the question. Processes of feedback, repair and expansion, for example, are central to Clark et al.'s view of conversation (Clark, 1996; Clark et al., 1983; Clark & Wilkes-Gibbs, 1986).

According to Clark and colleagues, not only do interlocutors collaborate during conversation, they also aim to minimize their joint effort:

Principle of Least Collaborative Effort

Participants in a contribution try to minimize the total effort spent on that contribution—in both the presentation and acceptance phases. Clark and Schaefer (1989, p. 269)

This principle was derived from observing the behavior of participants in a game of *Tangrams* or Chinese puzzles (Clark & Wilkes-Gibbs, 1986), as they decreased the length of their referring expressions and the number of turns over trials.

The view that redundant modification can be efficient if it facilitates the listener's visual search for the referent is in line with Clark et al.'s account of reference as a collaborative process (Rubio-Fernandez, 2016). However, in this account, collaboration need not involve turn taking, but a mutual understanding of the *joint goal of reference*. Thus, when a speaker requests a target in a referential communication task, their goal is to get the listener to identify the target in the array and, if the listener complies with the request, both interlocutors come to have the same goal. Given this joint goal, an optimal referential expression should be one that allows the listener to identify the intended target in the most efficient way; that is, with minimal expenditure of time and effort.

According to this view, it is an empirical question whether the shortest referential expression would be the most efficient one in any given context, as it has normally been assumed in the literature. Instead, redundant referential expressions should be understood as more or less efficient in a context rather than pragmatically infelicitous for being over-informative (Engelhardt et al., 2006, 2011). If we understand visual identification as part of the process of reference resolution (at least in situated reference), a pragmatic analysis should recognize that it is more efficient to refer to “the blue star” when the star is the only blue shape (as in the polychrome displays in Fig. 1) than when all shapes are blue

Fig. 1. Sample displays used to elicit spoken instructions (e.g., “Click on the blue star”) in Experiment 1a (Poly/2, Poly/4, Poly/8, Pop-Out/8 and Mono/4), Experiment 1b (Poly/8 and Pop-Out/8), Experiment 2 (Poly/4), and Experiment 3 (Size/2 and Size/16). The target in these displays was always the star. The participant grid corresponds to the 8-item conditions.

(as in the monochrome displays), regardless of the number of stars. Because standard pragmatic analyses are based on informativity alone, they distinguish underinformative, minimal, and overinformative expressions, but fail to account for the relative efficiency of redundant modification for visual search.

An efficiency-based account of referential communication also subsumes the traditional Gricean account based on informativity. Thus, an underinformative referential expression is less efficient than a minimal referential expression that establishes unique reference, insofar as the former leaves the listener either to choose randomly among various possible referents, or to ask for clarification. The advantage of an efficiency-based account is that discriminability may explain why cooperative speakers use redundant referential expressions in certain situations (see Rubio-Fernandez, 2016). By contrast, current pragmatic accounts can only treat redundancy as a departure from the Gricean Maxim of Quantity (e.g., Davies & Katsos, 2010; Engelhardt et al., 2006, 2011; Engelhardt & Ferreira, 2014; Grodner & Sedivy, 2011; Heller et al., 2012; Pogue et al., 2016; Sedivy, 2003, 2004, 2007).

In summary, an account of referential communication based on efficiency has greater explanatory power than standard pragmatic accounts based on informativity alone. The account I am developing here also responds to Sedivy’s plea for finer-grained accounts of referential communication: “Clearly, a simple notion of information that is required for the purpose of uniquely identifying a referent from among alternative referents is not sufficient” (2004, p. 361). By considering not only the amount of information provided by the speaker in relation to the number of competitors, but also how useful the description is for visual search, it may be possible to understand why people often use redundant modification, rather than simply treating those uses as violations of the Gricean Maxim of Quantity (1975).

4. Dividing the labor of reference: Between Grice and Clark

Given the above definition, the relative efficiency of a referential expression would be determined by the time and effort it takes a listener to identify the intended referent. However, the Gricean Maxims, including the Maxim of Quantity, were formulated for the speaker. That is, it is the speaker who must determine what is the minimal information they need to provide in order to ensure unique reference. Such calculations can be performed by keeping track of the competitors that the intended referent might have in the context (e.g., whether there are any other stars on the screen, apart from the blue one). However, if speakers are aiming to produce efficient referential expressions, rather than minimally informative ones, making an informed choice would ideally require knowing which referential expressions would result in faster identification times in a given context. Response time data, however, are not available to speakers wanting to produce an efficient referential expression.

What is available to a speaker formulating an object request is visual information about the environment in which the listener must identify the intended object. It is therefore on the basis of their visual common ground that a speaker must determine whether a redundant adjective may facilitate the listener's visual search for the referent. If speakers use redundant adjectives when it could be efficient (e.g., in polychrome displays), and refrain from doing so when it would not be (e.g., in monochrome displays), then that would suggest that redundant speakers are being cooperative. That is, even if they may not be minimally informative, as long as speakers are sensitive to the demands of the listener's visual search, their behavior would satisfy the two key assumptions in the Gricean program (1975); namely, that speakers are *rational* and *cooperative*. Therefore, an efficiency account of referential overspecification is in line with Gricean pragmatics, offering a more fine-grained analysis of the perceptual factors that a speaker may rely on in order to produce an optimal referential expression.

Also in line with Clark's (1996) view of language use, a speaker's effort to produce an optimal referential expression must take into account the listener's needs. Producing a modified description to pre-empt an ambiguity between two competitors is a canonical example of audience design: While the speaker may know which star they want, if there are two stars in the display, the speaker must compare them and specify which one they want for the listener's sake. Thus, in a context where there are two stars, the extra time and effort invested in producing and processing a color word is offset by the need to establish unique reference. However, in a context with a single star, using a redundant color adjective would only be justified if it facilitated the referent's identification relative to a minimal description. A number of empirical studies have suggested that redundant color adjectives can indeed facilitate the listener's search for a referent (Arts et al., 2011b; Mangold & Pobel, 1988; Paraboni & van Deemter, 2014; Paraboni et al., 2007; Rubio-Fernandez, 2017; Sonnenschein & Whitehurst, 1982; Tourtouri et al., 2019).

In situations of co-presence, however, producing a redundant description that facilitates the listener's visual search might not be so costly since the speaker can rely on mutual salience. That is, if the color of a referent seems salient to the speaker, they can assume

it will be salient for the listener too. In this vein, I want to propose that situations of co-presence allow speakers to rely on *visual heuristics* that can inform their choice of referential expression at a low processing cost (for other common ground heuristics, see Clark & Marshall, 1981). For example, the understanding that color is more efficient in polychrome than monochrome displays could be applied as a visual heuristic triggered by the perception of color contrast. Likewise, appreciating that the listener's visual search for a referent is harder the denser the display could be another visual heuristic that speakers can exploit in order to provide extra information at a low processing cost.

This type of low-cost pragmatics is in line with the distinction between *feedforward* and *recurrent processing audience design* drawn by Ferreira (2019), according to which speakers do not always engage in audience design through deliberative processes. In the case of feedforward audience design, speakers are able to make use of contextual features prior to the onset of utterance production and rely on previously learned strategies that facilitate communication. Recurrent processing audience design, by contrast, requires deliberative processes (e.g., in tasks where speaker and listener have different visual perspectives; e.g., Long et al., 2018; Wardlow-Lane & Ferreira, 2008).

The aim of this study was to test the hypothesis that overinformative speakers can be cooperative by relying on visual heuristics that help them produce efficient referential expressions. This view is in line with both Gricean pragmatics (1975) and Clark's (1996) view of reference as a collaborative process.

5. A study of cooperativeness in reference production

Following Rubio-Fernandez (2016), two types of factors affecting referential efficiency were considered in this study: perceptual factors and linguistic factors. Experiment 1a investigated three perceptual factors that might affect the relative efficiency of a color adjective for visual search: (i) color distinctiveness, (ii) density of the display, and (iii) perceptual salience. These three perceptual factors could be used as heuristics: (a) the higher the discriminatory power of a color adjective (with distinctive color being the extreme case), (b) the denser the display where the listener must identify the target, and (c) the more salient the color of the target in the display, the more efficient the use of a color adjective would be.

Color distinctiveness was tested in monochrome and polychrome displays as extreme cases where color adjectives would be either non-discriminating and therefore inefficient (Mono/4), or distinctive and potentially efficient (Poly/4). There is, of course, a wider range of visual displays where color may have more or less discriminatory power, and ongoing work is investigating speakers' sensitivity to more subtle visual manipulations. Likewise, density of the display was investigated here using polychrome displays of two, four, and eight shapes (Poly/2, Poly/4 and Poly/8), but there are obviously other combinations that could be used for this purpose (for earlier studies looking at the effect of density of the display on modification, see Clarke et al., 2013; Gatt et al., 2017; Koolen et al., 2015; Paraboni et al., 2007). Finally, visual salience was manipulated in pop-out

displays where all the shapes are of the same color, except the target (Pop-out/8). Pop-out displays included eight shapes in order to create a sharper pop-out effect that would be directly comparable to the densest polychrome condition. Pop-out displays have been extensively used to investigate efficient visual search, which takes place when the target has a distinctive property that makes it stand out from the distractors (Treisman & Gelade, 1980). Pop-out displays were used in the present study to investigate whether such visual effects also affect adjective production.

A further test of the cooperativeness hypothesis was performed by modifying the extent of the visual common ground shared by the interlocutors, rather than the visual displays. If the cooperativeness hypothesis is correct, speakers should produce redundant color adjectives when the listener needs to identify the referent in a display and color would guide their search, but not when the referent is part of their common ground. Thus, it was predicted that participants in Experiment 1a would produce more redundant color adjectives in the Poly/4 condition when their interlocutor did not know which shape was the target (i.e., as in a standard referential communication task) than when their interlocutor knew which shape was the target (i.e., as in a standard naming task). To further test the difference between these two types of task, the naming paradigm was used again in Experiment 1b with the visual materials from the Poly/8 and Pop-Out/8 conditions, in which color was most efficient/salient. Finally, Experiment 3 explored the density and salience heuristics in the production of redundant size adjectives, comparing participants' mention of size in displays of 2 shapes of different sizes and pop-out displays of 16 shapes where the target was of a different size than all the other shapes (Size/2 and Size/16).

With regard to linguistic factors, Experiment 2 examined the role of incrementality on the production of redundant modification. Because language is processed incrementally, redundant color adjectives are a more efficient visual cue in prenominal position because they can guide the listener's search for the referent (Eberhard et al., 1995; Spivey et al., 2001). Thus, an English speaker processing the instruction "Click on the blue star" would search for the referent by color, whereas a Spanish speaker processing the instruction "Haz clic en la estrella azul" would search for the same referent by shape, making the post-nominal adjective less efficient (for eye-tracking evidence, see Rubio-Fernandez et al., 2018). Thus, Experiment 2 investigated the use of redundant color adjectives in Poly/4 displays by English and Spanish speakers.

6. Experiment 1a

6.1. Method

6.1.1. Participants

A group of 120 students from University College London took part in the experiment. They were native speakers of English and participated for monetary compensation. All participants reported having normal color vision.

6.1.2. *Materials*

Two types of displays were created, one for the Experimenter (including 10 displays of geometrical shapes) and another one for the participants (including 10 empty grids with a cross indicating the position of the target in the Experimenter's display). The target shapes were as follows: arrow, circle, cylinder, heart, oval, pentagon, rectangle, square, star, and triangle, and they came in the following colors: blue (first and last trial), brown, green, gray, orange, pink, purple, red, and yellow. The number of shapes and colors in the displays varied, resulting in five different conditions (Mono/4, Poly/2, Poly/4, Poly/8 and Pop-Out/8; see Fig. 1 for sample items). The Pop-Out/8 condition was the only condition where participants could start anticipating the singleton target and, as a result, stop overspecifying color. For this reason, the Pop-Out/8 condition included five extra filler trials in which the target was a majority-color shape. Target position was counterbalanced across trials in all conditions, and trials were presented in the same random order to all participants.

The number of trials was kept low (10 critical trials in each condition) because previous studies had revealed that participants are highly consistent in their use of color adjectives across trials (Rubio-Fernandez, 2016; see also Tarenskeen et al. (2015) for an experimental investigation of response consistency in referential communication tasks).

6.1.3. *Procedure*

Two separate procedures were used in the task. In the referential communication task, participants sat beside and behind the Experimenter and had to ask the Experimenter to click on the target shape in each trial. Participants' requests were recorded and later coded as redundant or not redundant. Only referential expressions including both an adjective and a noun (e.g., "The blue star") were coded as redundant. Given the simplicity of the task, in order to avoid unnatural responses, participants were told that the original task was designed for children and that their responses would serve as control data. Participants were also asked to refer to the target object rather than to its position. The instructions stressed that all shapes in the displays were different (which implied that color was redundant in all trials) and that the Experimenter did not know which shape was the target in each trial. The task lasted less than 10 min.

A naming task was also used to validate the interpretation of the results of the referential communication task as evidence of cooperative behavior. Thus, in the naming task, the materials from the Poly/4 condition were presented on a computer monitor that the participant and the Experimenter shared as they sat side by side. In order to make the target shape part of their common ground, it was presented inside a square frame and participants were asked to name the shape that appeared inside a box in each display. The Experimenter moved forward the displays by pressing Enter.

Based on a previous study by Rubio-Fernandez (2016) using similar referential communication tasks, it was established a priori that 20 participants would be tested in each experimental condition. Conditions were manipulated between participants in order to avoid carry-over effects across different trial types (e.g., modification in a monochrome trial following modification in a polychrome trial), which would compromise the effect of

visual display. Previous referential communication studies have shown that manipulating the number of objects and colors within participants can lead to up to 60% of redundant modification in monochrome displays (Belke, 2006; Koolen et al., 2013) compared to 0% when testing the same conditions between participants (Rubio-Fernandez, 2016).

6.2. Results and discussion

The large majority of responses included a noun for the target shape, with only 2% of responses including a pronominal construction (e.g., “the blue one”), which were coded as non-redundant. No trial order effects were observed, with most participants adopting either a modification or a no-modification strategy across trials. The mean proportions of redundant color adjectives observed in Experiment 1a are reported in Fig. 2. The results of the four experiments reported in this study were analysed using non-parametric statistics because the data were not normally distributed, and the extreme values observed across conditions interfered with model convergence when mixed model analyses were attempted.

Looking first at the effect of color distinctiveness, participants produced more redundant color adjectives in Poly/4 than in Mono/4, in which they never mentioned the color of the target. A Mann-Whitney test carried out on the proportion of redundant color adjectives revealed a significant difference between the two conditions ($U = 290$; $Z = -2.42$, $p = .016$).

Regarding display density, a Kruskal–Wallis test carried out on the proportions of redundant color adjectives elicited in the Poly/2, Poly/4, and Poly/8 conditions revealed a significant effect ($H(2) = 6.02$, $p = .049$) in the direction predicted by the cooperativeness hypothesis: Speakers produced more color adjectives the denser the display. Mann–Whitney tests comparing adjective production in the Poly/2 and Poly/4 conditions did not

Fig. 2. Mean proportions of redundant color and minimal descriptions observed across the different conditions of Experiment 1a.

reveal a significant difference ($U = 234$; $Z = 0.89$, $p = .374$), and the same was observed between Poly/4 and Poly/8 ($U = 153$; $Z = 1.27$, $p = .204$), but a significant difference was observed between Poly/2 and Poly/8, ($U = 295$; $Z = 2.54$, $p = .011$).

Regarding the effect of color salience, a Mann–Whitney test comparing the proportions of color adjectives observed in Poly/8 and Pop-Out/8 revealed a marginally significant advantage in the Pop-Out/8 condition ($U = 265$; $Z = 1.74$, $p = .082$), supporting the cooperativeness hypothesis.

Finally, participants did not produce redundant color adjectives when asked to name the target shapes in the Poly/4 condition. A Mann–Whitney test revealed a significant difference between the proportion of color adjectives observed in the two Poly/4 conditions ($U = 290$; $Z = -2.42$, $p = .016$). This pattern of results supports the hypothesis that participants in the referential communication task produced redundant color adjectives because the Experimenter had to search for the referent, not doing so when the referent was known to both interlocutors.

Overall, the results of Experiment 1a support the prediction that redundant speakers are cooperative, using color adjectives more often when it would more likely facilitate the listener's visual search for the referent. Moreover, participants revealed high sensitivity to the visual context, rather than a general strategy to use color in polychrome displays.

7. Experiment 1b

The rates of redundant modification observed in the Poly/4 condition were significantly lower when participants were naming the shapes than when they were referring to a target for the benefit of the Experimenter. However, it is possible that the difference between these two tasks would not be so stark if participants were naming shapes in Poly/8 or Pop-Out/8 displays, since those elicited the highest rates of redundant modification in Experiment 1a. Such a pattern of results would suggest that color salience may trigger the use of redundant color adjectives, regardless of the listener's task.

7.1. Method

7.1.1. Participants

A group of 30 students from the Massachusetts Institute of Technology took part in the study. They were native speakers of English and participated for monetary compensation. All participants reported having normal color vision.

7.1.2. Materials and procedure

The visual displays from conditions Poly/8 and Pop-Out/8 in Experiment 1a were combined in a random sequence in Experiment 1b, starting with a Pop-Out/8 display. These two conditions were tested together because they had elicited the highest rates of

redundant modification in the first experiment and carryover effects would run counter to the cooperativeness hypothesis, thus providing a more stringent test. The procedure was the same that was used in Experiment 1a for the naming task.

7.2. Results

Of the 30 participants in the task, only one used redundant color adjectives in all trials. Another participant produced redundant modification in 28% of trials. None of the remaining participants produced color adjectives, resulting in an average rate of redundant modification of 4%. This is markedly below the 75% observed in Experiment 1a for conditions Poly/8 and Pop-Out/8 combined (see Fig. 3).

I have argued that in situations of co-presence, it is not possible to establish whether a speaker is using a redundant color adjective because the color of the referent seemed salient to them, or because they assumed it would also be salient for the listener. However, what the starkly different results of the referential communication and naming tasks suggest is that even when the color of a target might be most salient to the speaker (e.g., in a pop-out display), speakers do not produce redundant color adjectives if the listener need not search for the target in the display. This means that even in situations of co-presence where the speaker can rely on their own perspective, speakers are sensitive to their listener’s needs—thus behaving cooperatively.

Fig. 3. Mean proportions of redundant color and minimal descriptions observed in the Poly/8 and Pop-Out/8 conditions of Experiment 1a (combined across participants) and Experiment 1b (combined within participants). The two conditions were used in a referential communication task in the first experiment, and in a naming task in the second experiment.

8. Experiment 2

The third experiment investigated whether speakers appreciate that redundant color adjectives are more efficient in prenominal position than in postnominal position. Rubio-Fernandez (2016) already tested this hypothesis with positive results. However, the participants in that study had to refer to items of clothing, which elicited high rates of redundant color adjectives across conditions (probably because color is a central property of clothes). Experiment 2 tried to establish whether English speakers would also use redundant color adjectives more often than Spanish speakers when referring to geometrical shapes (which are not so closely related to color).

This experiment also tried to control for an alternative interpretation of the results of Rubio-Fernandez (2016): Because color adjectives are encoded in postnominal position, Spanish speakers have more time to decide whether color is necessary to establish unique reference. Therefore, Spanish speakers may simply be more accurate in their use of color modification, rather than the difference with English speakers being related to the relative efficiency of prenominal and postnominal adjectives. In order to control for this possibility, the presentation of the displays was timed in Experiment 2 so that both groups of participants had the same time to formulate their requests.

8.1. Method

8.1.1. Participants

In this study, 20 students from Princeton University and 20 students from the University of Oviedo (Spain) took part in the experiment. The Princeton students were native speakers of English and the Oviedo students were native speakers of Spanish. Both groups participated for monetary compensation. All participants reported having normal color vision.

8.1.2. Materials and procedure

The materials from the Poly/4 condition in Experiment 1a were used again in Experiment 2. The procedure was also the same with the only difference that the presentation of the displays was timed for 1.5 s. After that time, the display disappeared from the monitor. The alleged explanation for this was that the task was otherwise too easy for adult participants since it was originally designed for children.

8.2. Results and discussion

No trial order effects were observed, with most participants adopting either a modification or a no-modification strategy across trials. The mean proportions of redundant color adjectives produced by English and Spanish speakers are reported in Fig. 4. A Mann–Whitney test was carried out on participants' proportions of color modification revealing a significant difference between the two language groups ($U = 124$; $Z = 2.06$, $p = .039$).

The results of Experiment 2 replicate the findings by Rubio-Fernandez (2016) but using geometrical shapes, which overall elicited lower rates of color overspecification. A significant difference between English and Spanish speakers was observed even when the presentation of the displays was timed and both groups had comparable time to produce their instructions. This manipulation rules out the possibility that Spanish speakers simply had more time to plan their descriptions. Therefore, the results of Experiment 2 support the hypothesis that speakers are sensitive to linguistic factors affecting the relative efficiency of redundant modification, producing color adjectives to different degrees in languages with different adjective positions.

9. Experiment 3

The results of Experiments 1, 2, and 3 supported the hypothesis that redundant speakers can be cooperative. However, this hypothesis was only tested with color adjectives, leaving open the possibility that these results are related to the salience of color as a visual property. Moreover, previous studies have shown that size adjectives are used redundantly less often than color adjectives (e.g., Belke & Meyer, 2002; Brown-Schmidt & Konopka, 2008; Engelhardt & Ferreira, 2014; Pechmann, 1989; Tarenskeen et al.,

Fig. 4. Mean proportions of redundant color and minimal descriptions observed in the two language groups in Experiment 2.

2015). Therefore, the last experiment of the study tested the cooperativeness hypothesis with size adjectives.

Participants in Experiment 3 had to refer to a target shape that was of a different size than the other shape(s) in either a 2-shape display or a 16-shape display (see Fig. 1). Redundant modification would be more efficient in the 16-shape condition than in the 2-shape condition because in the former, the listener would be performing a harder visual search than in the latter, benefitting more from a redundant size adjective. Relatedly, participants in this task could rely on two visual heuristics: the density of the display and the salience of the target's size. The density of the display could serve as a proxy for the relative difficulty of the listener's visual search. Also, because of the pop-out effect, the odd size of the target would be highly salient in the denser displays, both for the speaker and the listener.

9.1. Method

9.1.1. Participants

A set of 40 students from University College London participated in the experiment. They were all native speakers of English and participated for monetary compensation.

9.1.2. Materials and procedure

Two sets of materials were built, one for each condition. For the Size/2 condition, 10 displays were constructed including two different shapes of different sizes. The target was the smaller shape in half the trials. The target shapes were the same as in previous experiments. All the displays were monochrome to avoid that color modification may trigger size modification.

The two shapes used in the Size/2 displays were used again in the Size/16 condition but inserted into a 4×4 display, thereby keeping the relative size difference constant (see Fig. 1). Dense displays of 16 shapes were used in order to create a sharp pop-out effect that could elicit redundant size modification. In each 4×4 display, there was an odd-size shape that was either smaller or larger than all the other shapes and which corresponded, in critical trials, with the target of the Size/2 condition. In order to avoid the singleton target becoming predictable, 10 filler trials were constructed in which the target shape was of the same size as the majority (and larger than the singleton in half the trials). The position of the odd-size shape was counterbalanced across trials. Trials were presented in the same random order to all participants.

As in previous experiments, participants had to ask the Experimenter to click on a target shape in a computer display. The procedure was the same as in Experiments 1a and 3, with participants being given printouts of empty grids with a cross indicating the position of the target in the Experimenter's display. Again, participants were told that their responses would serve as control data in a study originally designed for children.

9.2. Results and discussion

No trial order effects were observed, with most participants adopting either a modification or a no-modification strategy across trials. The mean proportions of redundant size adjectives observed in Experiment 3 are reported in Fig. 5. Participants did not overspecify the size of the target in any of the 2-shape displays, but did so 61% of the time when size was distinctive in the 16-shape displays. In the filler trials of the Size/16 condition, in which size was not distinctive, participants produced redundant modification 15% of the time. It must be noted, however, that in 16 of the 24 filler trials in which participants produced a redundant size adjective, there could have been a spill-over effect from the previous trial, which was a critical one.

A Mann–Whitney test was carried out on the proportions of redundant size adjectives, revealing a significant difference between the Size/2 and Size/16 conditions ($U = 350$; $Z = -4.04$, $p < .001$). The results of Experiment 3 therefore confirm that speakers produce not only redundant color adjectives but also redundant size adjectives in those visual contexts where it would have more discriminatory power and be more beneficial for their interlocutor. These results are also in line with previous studies showing that people tend to produce longer (and often redundant) expressions to refer to entities in visually cluttered environments relative to their referential expressions in more sparse displays (e.g., Clarke et al., 2013; Gatt et al., 2017; Koolen et al., 2015; Paraboni et al., 2007).

10. General discussion

Standard pragmatic accounts in the Gricean tradition establish referential specification in relation to a referent's competitors in the situational context (e.g., Davies & Katsos, 2010; Engelhardt et al., 2006, 2011; Heller et al., 2012; Pogue et al., 2016; Sedivy, 2003, 2004): A modified referential expression is pragmatically felicitous if there are other objects of the same category as the intended referent and the adjective pre-empts an ambiguity; otherwise, the same expression is considered overinformative. In this paper, I have developed the alternative pragmatic account of referential specification first proposed by Rubio-Fernandez (2016), which takes into account not only informativity, but also discriminability and incrementality of processing. According to this account, in order to determine the relative efficiency of a given referential expression, one must establish not only the number of competitors in the context, but also how the referent differs from all the other objects in terms of color, size, and other properties that could make the referent visually salient. Thus, referring to a star as "the blue star" in a situation where there is only one star would be redundant according to the standard pragmatic analysis. However, the same expression could be optimal according to an efficiency-based account if the star was the only blue shape among various shapes of different colors (Rubio-Fernandez, 2016).

This pragmatic account allows establishing the relative efficiency of speaker-listener coordination in referential communication. Rather than assuming that referential overspecification is either "for the speaker" or "for the listener" (as it is often debated in the

Fig. 5. Mean proportions of redundant size and minimal descriptions observed in the two visual conditions in Experiment 3.

literature; for a recent review, see Davies & Arnold, 2019), I started by pointing out that situations of co-presence (such as those recreated in standard referential communication tasks) are not ideal for the investigation of perspective taking because speaker and listener share the same perspective. Instead, referential communication tasks such as the ones used in this study allow investigating whether speakers use potentially efficient referential expressions that could facilitate the listener's search for the referent, without incurring in perspective-taking costs. Thus, in situations of co-presence, speakers can rely on visual heuristics, such as the perception of color contrast, the relative density of the display, or the visual salience of the referent, which facilitate the formulation of referential expressions that would be efficient for the listener. This account is in line with Ferreira's feedforward audience design (2019), which employs learned communicative strategies to accommodate an interlocutor's needs with minimal expenditure of executive resources.

In four language-production experiments, I tested the hypothesis that speakers are cooperative and produce redundant modification when it may facilitate the listener's visual search for the referent (see also Rubio-Fernandez, 2016). The results offered support to this hypothesis, with speakers producing redundant color adjectives when (a) color was a distinctive property of the referent, (b) the display was denser with shapes, and (c)

the color of the target was salient in the display (Experiment 1a). In addition, a control condition using the same visual materials in a naming task revealed that people do not produce redundant color adjectives when the intended referent is part of their common ground with their interlocutor, supporting the view that when they do so, they are being sensitive to the listener's visual search for the referent (Experiments 1a and 1b).

The next experiment investigated referential efficiency in relation to incrementality. Prenominal but not postnominal adjectives, can guide a listener's visual search for a referent, making redundant color adjectives more efficient in prenominal than in postnominal position. Supporting the cooperativeness hypothesis, English speakers produced more redundant color adjectives than Spanish speakers (Experiment 2). Finally, speakers produced redundant size adjectives in denser displays where the size of the target was highly salient (Experiment 3), extending the results with color adjectives to size adjectives. Overall, the results of these language-production experiments support the view that speakers can be cooperative when they produce adjectives that are not, strictly speaking, necessary.

It must be noted that the results of this study may not extrapolate to other speech acts where the interlocutor need not find the referent in the shared physical environment (see, e.g., Yoon et al., 2012). In addition, efficiency is not the only factor affecting the production of redundant adjectives in object requests. Rubio-Fernandez (2016; Experiment 1) observed that people often used redundant color adjectives when asking for clothes in monochrome displays (where color would be useless as a visual cue), in contrast with what was observed in the present study when participants had to refer to geometrical shapes in monochrome displays (Experiment 1a). I explain this difference as an effect of *color pertinence*, which varies depending on the semantic category of the noun: While color is a central property of clothes and fashion items (as suggested by collocations such as "black tie" or "little black dress"), it is not so central for other semantic categories, such as geometrical shapes. The effect of semantic category on the production of color adjectives is another linguistic factor that needs to be considered in modelling referential specification.

Rubio-Fernandez (2016, Experiment 2) also observed that when participants had to refer to fruits and vegetables of predictable colors (e.g., a banana or a carrot), they did not mention their color, even if it was discriminatory in the display (see also Sedivy, 2003; Westerbeek et al., 2015). These patterns of results suggest that color pertinence and color predictability for a given noun affect the production of redundant color adjectives above and beyond considerations of discriminability and visual salience, suggesting that linguistic factors can weigh more heavily than perceptual factors in the production of redundant modification. However, while an efficiency-based account of referential communication does not explain all uses of redundant modification, it has more explanatory power than standard pragmatic accounts based on informativity alone. Such traditional accounts cannot explain the variability observed in the language-production experiments reported in this paper, nor the potential advantage of redundant modification for visual search.

Acknowledgments

Thanks to Bart Geurts for many interesting discussions of color, efficiency, and cooperativity over the years. Thanks also to Madeleine Long for assistance with data collection. This research was supported by a *Young Research Talent Grant* from the Research Council of Norway (*Linguistic and cognitive factors in effective referential communication*; Ref. 230718) and by a *Euro-XPrag Grant* from the European Science Foundation (*The cognitive costs and benefits of using color adjectives*; Ref. 3844). The author gratefully acknowledges this funding. Final thanks to five anonymous reviewers for their constructive comments and insightful suggestions.

References

- Arnold, J. (2008). Reference production: Production-internal and addressee-oriented processes. *Language and Cognitive Processes*, 23, 495–527.
- Arts, A., Maes, A., Noordman, L., & Jansen, C. (2011a). Overspecification in written instruction. *Linguistics*, 49, 555–574.
- Arts, A., Maes, A., Noordman, L., & Jansen, C. (2011b). Overspecification facilitates object identification. *Journal of Pragmatics*, 43, 361–374.
- Barr, D. J. (2014). Perspective taking and its impostors in language use: Four patterns of deception. In T. Holtgraves (Ed.), *The Oxford handbook of language and social psychology* (pp. 98–110). Oxford, UK: Oxford University Press.
- Belke, E. (2006). Visual determinants of preferred adjective order. *Visual Cognition*, 14, 261–294.
- Belke, E., & Meyer, A. S. (2002). Tracking the time course of multidimensional stimulus discrimination: Analyses of viewing patterns and processing times during “same”–“different” decisions. *European Journal of Cognitive Psychology*, 14, 237–266.
- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 22(6), 1482–1493.
- Brennan, S. E., Galati, A., & Kuhlén, A. K. (2010). Two minds, one dialog: Coordinating speaking and understanding. *Psychology of Learning and Motivation*, 53, 301–344.
- Brown-Schmidt, S., & Konopka, A. E. (2008). Little houses and casas pequeñas: Message formulation and syntactic form in unscripted speech with speakers of English and Spanish. *Cognition*, 109, 274–280.
- Clark, H. H. (1996). *Using language*. Cambridge, UK: Cambridge University Press.
- Clark, H. H., & Marshall, C. R. (1981). Definite reference and mutual knowledge. In A. K. Joshi, B. Webber, & I. A. Sag (Eds.), *Elements of discourse understanding* (pp. 10–63). Cambridge: Cambridge University Press.
- Clark, H. H., & Schaefer, E. F. (1989). Contributing to discourse. *Cognitive Science*, 13, 259–294.
- Clark, H. H., Schreuder, R., & Buttrick, S. (1983). Common ground at the understanding of demonstrative reference. *Journal of Verbal Learning and Verbal Behavior*, 22(2), 245–258.
- Clark, H. H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22, 1–39.
- Clarke, A. D., Elsner, M., & Rohde, H. (2013). Where’s Wally: The influence of visual salience on referring expression generation. *Frontiers in Psychology*, 4, 329.
- Davies, C., & Arnold, J. E. (2019). Reference and informativeness: How context shapes referential choice. In C. Cummins & N. Katsos (Eds.), *Handbook of experimental semantics and pragmatics* (pp. 474–493). Oxford, UK: Oxford University Press.
- Davies, C., & Katsos, N. (2010). Over-informative children: Production/comprehension asymmetry or tolerance to pragmatic violations? *Lingua*, 120, 1956–1972.

- Eberhard, K. M., Spivey-Knowlton, M. J., Sedivy, J. C., & Tanenhaus, M. K. (1995). Eye movements as a window into real-time spoken language comprehension in natural contexts. *Journal of Psycholinguistic Research*, 24(6), 409–436.
- Engelhardt, P. E., Bailey, K. G. D., & Ferreira, F. (2006). Do speakers and listeners observe the Gricean Maxim of Quantity? *Journal of Memory and Language*, 54, 554–573.
- Engelhardt, P. E., Demiral, Ş. B., & Ferreira, F. (2011). Over-specified referring expressions impair comprehension: An ERP study. *Brain and Cognition*, 77, 304–314.
- Engelhardt, P. E., & Ferreira, F. (2014). Do speakers articulate over-described modifiers differently from modifiers that are required by context? Implications for models of reference production. *Language and Cognitive Processes*, 29, 975–985.
- Engelhardt, P. E., & Ferreira, F. (2016). Reaching sentence and reference meaning. P. Knoeferle, P. Pykkönen-Klauck, & M. W. Crocker (Eds.). *Visually Situated Language Comprehension*, 93, 127.
- Ferreria, V. S. (2019). A mechanistic framework for explaining audience design in language production. *Annual Review of Psychology*, 70, 29–51.
- Ford, W., & Olson, D. (1975). The elaboration of the noun phrase in children's description of objects. *Journal of Experimental Child Psychology*, 19(3), 371–382.
- Gatt, A., Krahmer, E., Van Deemter, K., & van Gompel, R. P. (2017). Reference production as search: The impact of domain size on the production of distinguishing descriptions. *Cognitive Science*, 41, 1457–1492.
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Speech acts* (pp. 41–58). New York: Academic Press.
- Grodner, D., & Sedivy, J. (2011). The effect of speaker-specific information on pragmatic inferences. In N. Pearlmuter & E. Gibson (Eds.), *The processing and acquisition of reference* (pp. 239–272). Cambridge, MA: MIT Press.
- Heller, D., Gorman, K. S., & Tanenhaus, M. K. (2012). To name or to describe: Shared knowledge affects referential form. *Topics in Cognitive Science*, 4, 290–305.
- Horton, W. S., & Brennan, S. E. (2016). The role of metarepresentation in the production and resolution of referring expressions. *Frontiers in Psychology*, 7. Special issue *Models of Reference*.
- Keysar, B. (1997). Unconfounding common ground. *Discourse Processes*, 24, 253–270.
- 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.
- Koolen, R., Krahmer, E., & Swerts, M. (2015). How distractor objects trigger referential overspecification: Testing the effects of visual clutter and distractor distance. *Cognitive Science*, 40, 1617–1647.
- Long, M. R., Horton, W. S., Rohde, H., & Sorace, A. (2018). Individual differences in switching and inhibition predict perspective-taking across the lifespan. *Cognition*, 170, 25–30.
- Maes, A., Arts, A., & Noordman, L. (2004). Reference management in instructive discourse. *Discourse Processes*, 37, 117–144.
- Mangold, R., & Pobel, R. (1988). Informativeness and instrumentality in referential communication. *Journal of Language and Social Psychology*, 7, 181–191.
- Olson, D. R. (1970). Language and thought: Aspects of a cognitive theory of semantics. *Psychological Review*, 77(4), 257–273.
- Paraboni, I., & van Deemter, K. (2014). Reference and the facilitation of search in spatial domains. *Language, Cognition and Neuroscience*, 29, 1002–1017.
- Paraboni, I., van Deemter, K., & Masthoff, J. (2007). Generating referring expressions: Making referents easy to identify. *Computational Linguistics*, 33, 229–254.
- Pechmann, T. (1989). Incremental speech production and referential overspecification. *Linguistics*, 27, 89–110.
- Pogue, A., Kurumada, C., & Tanenhaus, M. K. (2016). Talker-specific generalization of pragmatic inferences based on under- and over-informative pronominal adjective use. *Frontiers in Psychology*, 6. Special Issue *Models of Reference*.

- Rubio-Fernandez, P. (2017). Redundant color words are more efficient than shorter descriptions. Paper presented at the *CUNY Conference on Sentence Processing*. Massachusetts Institute of Technology.
- 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., Mollica, F., & Jara-Ettinger, J. (2018). English and Spanish speakers exploit language structure to increase communicative efficiency. Poster presented at the *CUNY Conference on Sentence Processing*. University of California at Davis. <https://doi.org/10.31234/osf.io/gf8qx>
- 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 & M. Tanenhaus (Eds.), *Approaches to studying world-situated language use* (pp. 345–364). Cambridge, MA: MIT Press.
- Sedivy, J. C. (2007). Implicature during real time conversation: A view from language processing research. *Philosophy Compass*, 2, 475–496.
- Sonnenschein, S., & Whitehurst, G. J. (1982). The effects of redundant communications on the behavior of listeners: Does a picture need a thousand words? *Journal of Psycholinguistic Research*, 11, 115–125.
- Spivey, M. J., Tyler, M. J., Eberhard, K. M., & Tanenhaus, M. K. (2001). Linguistically mediated visual search. *Psychological Science*, 12(4), 282–286.
- Tarenskeen, S., Broersma, M., & Geurts, B. (2015). Overspecification of color, pattern, and size: Salience, absoluteness, and consistency. *Frontiers in Psychology*, 6. Special Issue *Models of Reference*.
- Tourtour, E. N., Delogu, F., Sikos, L., & Crocker, M. W. (2019). Rational over-specification in visually-situated comprehension and production. *Journal of Cultural Cognitive Science*, 1–28.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- van der Sluis, I., & Krahmer, E. (2007). Generating multimodal references. *Discourse Processes*, 44, 145–174.
- Wardlow-Lane, W., & Ferreira, V. S. (2008). Speaker-external versus speaker-internal forces on utterance form: Do cognitive demands override threats to referential success? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 1466–1481.
- Westerbeek, H., Koolen, R., & Maes, A. (2015). Stored object knowledge and the production of referring expressions: The case of color typicality. *Frontiers in Psychology*, 7. Special issue *Models of Reference*.
- Yoon, S. O., Koh, S., & Brown-Schmidt, S. (2012). Influence of perspective and goals on reference production in conversation. *Psychonomic Bulletin & Review*, 19, 699–707.