



ARTICLE



<https://doi.org/10.1057/s41599-021-00818-6>

OPEN

Color discriminability makes over-specification efficient: Theoretical analysis and empirical evidence

Paula Rubio-Fernandez ^{1,2} 

A psychophysical analysis of referential communication establishes a causal link between a visual stimulus and a speaker's perception of this stimulus, and between the speaker's internal representation and their reference production. Here, I argue that, in addition to visual perception and language, social cognition plays an integral part in this complex process, as it enables successful speaker-listener coordination. This pragmatic analysis of referential communication tries to explain the redundant use of color adjectives. It is well documented that people use color words when it is not necessary to identify the referent; for instance, they may refer to "the blue star" in a display of shapes with a single star. This type of redundancy challenges influential work from cognitive science and philosophy of language, suggesting that human communication is fundamentally efficient. Here, I explain these seemingly contradictory findings by confirming the *visual efficiency hypothesis*: redundant color words can facilitate the listener's visual search for a referent, despite making the description unnecessarily long. Participants' eye movements revealed that they were faster to find "the blue star" than "the star" in a display of shapes with only one star. A language production experiment further revealed that speakers are highly sensitive to a target's discriminability, systematically reducing their use of redundant color adjectives as the color of the target became more pervasive in a display. It is concluded that a referential expression's efficiency should be based not only on its informational value, but also on its discriminatory value, which means that redundant color words can be more efficient than shorter descriptions.

¹Department of Philosophy, University of Oslo, Oslo, Norway. ²Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA. email: paula.rubio-fernandez@ifikk.uio.no

Introduction

Reference is one of the most fundamental functions of communication and therefore, humans have at their disposal a wide variety of referential forms: from an ostensive glance or a pointing gesture, to the use of pronouns (e.g., “I saw that”) or complex noun phrases (e.g., “I met the young couple who has a small dog”). Referents may also be abstract (e.g., “the main problem”) or physical (e.g., “the blue cup”), and may be co-present to both interlocutors (e.g., “Is that yours?”), present to only one of them (e.g., “I’ll keep them safe until you arrive”), or absent altogether (e.g., “Mia left this morning”). In addition, reference can be performed in the gestural, verbal or written modalities. Against this wide range of referential forms, referent types, communicative settings and modalities, referential communication studies such as the present one focus on a rather specific form of reference: verbal reference to physical objects in interactive settings (e.g., requesting “the blue cup” when sitting with others at a breakfast table). While limited in scope, this focus on physically grounded reference during interaction allows investigating three cognitive abilities that are often recruited in interactive referential communication, but which are not always included in theoretical and computational models of reference: namely, pragmatics, visual perception and social cognition. The present paper has two aims, one theoretical and one empirical, which will be addressed in turn. The first aim is to develop a pragmatic account of physical reference that integrates pragmatics, visual perception and social cognition into the use and comprehension of color adjectives. The second aim is to test a related hypothesis—the *visual efficiency hypothesis*, according to which interlocutors in referential communication are rational and cooperative as they jointly exploit the discriminability of color.

In the last decade, research in cognitive science has uncovered how human language is shaped by a pressure to communicate efficiently (Gibson et al., 2019). Yet contrary to this view, experimental studies have repeatedly shown that people often use adjectives, especially color, when it is not necessary to identify the referent (e.g., Sonnenschein and Whitehurst, 1982; Mangold and Pobel, 1988; Belke and Meyer, 2002; Sedivy, 2003, 2004; Maes et al., 2004; Belke, 2006; Paraboni et al., 2007; Paraboni and van Deemter, 2014). For example, someone may ask you “Pass me the blue cup” when there is only one cup on the table (i.e., when there are no *competitors* in the visual context), instead of using the shorter instruction “Pass me the cup.” This finding has generally been interpreted as a challenge to the standard theory of communication proposed by philosopher Paul Grice (1975), according to which speakers should not make their contributions more informative than necessary for the purpose of the exchange. In line with the Gricean Maxim of Quantity, the general assumption in the experimental literature is that shorter referential expressions are preferable over redundant ones because the latter require extra effort to produce and time to process, with no added informational value (e.g., Pechmann, 1989; Sedivy, 2003, 2004; Engelhardt et al., 2006, 2011; Koolen et al., 2013). This means color words should only be used when they are necessary (e.g., to distinguish two cups of different colors) and have informational value; otherwise shorter descriptions are preferred.

Two caveats are in order at this point:

Caveat 1: The widespread view that using color adjectives redundantly violates the Gricean Maxim of Quantity rests on an untested assumption. In order for an adjective to provide more information than is required “for the current purposes of the exchange” (Grice, 1975: p. 45), it would need to be used contrastively (or technically speaking, *restrictively*; Umbach, 2006) in the absence of competitors. For example, a speaker would use a color adjective contrastively if they requested “the blue cup” to distinguish it from a red cup. However, color adjectives may also

be used descriptively (or *non-restrictively*; e.g., “Today we visited the big red farm near the motorway”). Given this pragmatic distinction, a color adjective would only violate the Gricean Maxim of Quantity if the speaker used it contrastively when no modification was necessary. Such a situation would arise if a speaker requested “the blue cup” from a cluttered breakfast table, for example, on the mistaken assumption that there must be other cups on the table. Research on referential over-specification has often argued that indeed, speakers use color adjectives redundantly in order to pre-empt a possible ambiguity when they have not scanned the scene for competitors (Pechmann, 1989; Belke and Meyer, 2002; Belke, 2006; Engelhardt et al., 2006; Koolen et al., 2013; Fukumura and Carminati, 2021). However, to the best of my knowledge, this general assumption remains untested.¹

Caveat 2: The standard view of referential efficiency (i.e., say as much as necessary, but not more) is based solely on information quantity (measured in number of words) and fails to account for perceptual factors. Perception is tightly related to physically grounded reference: as a speech act, the goal of physical reference is to ensure that the listener identifies the intended referent in the visual context. From this standpoint, an efficient referential expression is one that achieves this goal fast and easily (Rubio-Fernandez, 2016, 2019). Providing sufficient information for the listener to identify a referent undoubtedly contributes to a referential expression’s efficiency, since otherwise the listener is left to guess or ask for clarification (e.g., “Which cup, the blue one or the red one?”). However, the relative ease or difficulty of the listener’s visual search must also count towards a referential expression’s efficiency. Thus, even if a so-called “redundant color word” does not have *informational value* (because it does not distinguish the target referent from other objects of the same kind), it may have *discriminatory value* if it facilitates the listener’s identification of the referent in the entire visual context (see Fig. 1, adapted from Rubio-Fernandez et al., 2020).

The present study will test the visual efficiency hypothesis, according to which speakers are rational and cooperative if they use redundant color adjectives when color has discriminatory

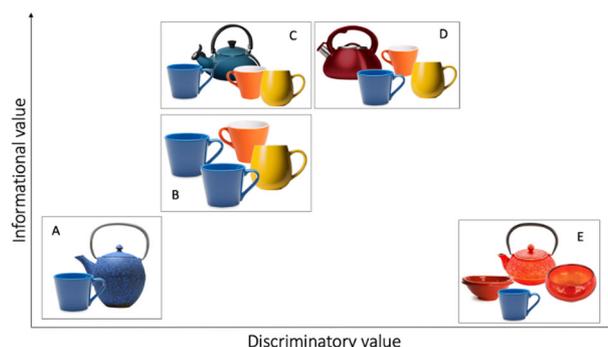


Fig. 1 Visualization of the informational vs. discriminatory value of the color adjective “blue” in the description “the blue cup”. In Displays **A** and **E**, “blue” has no informational value because there is only one cup (i.e., color does not pre-empt an ambiguity). In Display **B**, “blue” has more informational value because there are several cups, but not enough to secure unique reference amongst them. In Displays **C** and **D**, “blue” has the highest informational value because it resolves reference among all cups. Regarding discriminatory value, “blue” has none in Display **A** because both objects are blue, whereas it has more in Displays **B** and **C** because two of the four objects are blue. In Display **D**, “blue” has higher discriminatory value because there is a single blue object, and it has even more in Display **E** because it relies on a pop-out effect (i.e., the uniform color of the other objects makes the color blue stand out).

value in the visual context, despite lacking informational value in the absence of competitors. Contrary to the standard analysis, this view of modification does not assume that color adjectives are used restrictively in referential communication (i.e., with the aim of pre-empting a possible ambiguity with undetected competitors) but rather descriptively (i.e., with the aim of facilitating the listener's visual search for the referent). In this view, color contrast is established across the entire visual display (e.g., the blue object vs. all non-blue objects) rather than between competitors of the same kind (e.g., the blue cup vs. all other cups; for eye-tracking evidence, see Rubio-Fernandez et al., 2020; Rubio-Fernandez and Jara-Ettinger, 2020).²

Imagine a situation where there is a single cup on a breakfast table: under an informativity analysis, the description “the blue cup” would be equally redundant if the cup was the only blue object on the table, or if all the tableware were blue. Yet for visual search, the word “blue” is clearly more efficient when it is a distinctive property of the referent than when all objects are the same color (Rubio-Fernandez, 2016, 2019). More generally, an extensive literature on visual cognition has shown that color plays a key role in object recognition (see, e.g., Davidoff, 1991, 2001; Gegenfurtner and Rieger, 2000; Tanaka et al., 2001; Bramão et al., 2011; Adams and Chambers, 2012). In addition, color adjectives have a special status in reference production research, because color is an absolute property that is visually salient, unlike size, material or orientation (see Sedivy, 2003, 2004; Gatt et al., 2014). Therefore, it is ultimately an empirical question whether a redundant color word may facilitate the listener's visual search for a referent sufficiently to be more efficient than a shorter, unmodified description.

The present study is the first to address this empirical question by employing eye-tracking during real-time language processing. A second experiment looks at the effect of discriminability on the production of redundant color adjectives. Together, these two experiments investigate the role of discriminability in language comprehension and language production. However, before we turn to the empirical contribution of the study, I will develop a pragmatic analysis of referential communication that integrates visual perception and social cognition, rather than focusing solely on informativity. This new theoretical analysis will borrow insight from *psychophysics*—the branch of psychology that studies the relationship between a physical stimulus in the outside world and how its perception is connected to human performance.

The paper is divided in two parts. The first part develops a theoretical account of referential communication, which moves away from standard views on informativity (§2) and adopts instead a psychophysical perspective (§3). The role of social cognition in referential communication will be discussed next (§4), followed by a brief review of recent work on efficient audience design (§5). In the second part of the paper, I will discuss the effect of over-specification on visual search (§6) and report a referential communication study (§7), including an eye-tracking reference resolution task (§8) and a web-based reference production task (§9). The paper closes with a discussion of the implications of visual efficiency for research on color over-specification (§10).

Recent advances on informativity

Following Grice's seminal work (1975), pragmatic accounts of referential communication have focused on informativity (for a review, see Davies and Arnold, 2019), distinguishing amongst *under-informative expressions* (e.g., referring to “the blue cup” in a situation with several blue cups, as in Fig. 1B), *optimally informative expressions* (e.g., referring to “the blue cup” when there are several cups, only one of which is blue, as in Fig. 1C),

and *over-informative expressions* (e.g., referring to “the blue cup” when there is only one cup, as in Fig. 1E). The latest advances on the role of informativity in referential communication have been made in computational research. Degen et al. (2020) have recently put forward a computational model of reference production that relaxes the assumption that informativity is computed with respect to a deterministic semantics, proposing instead that “certain terms may apply better than others to an object without strictly being true or false” (p. 592; for an early precursor of this idea in terms of fuzzy logic, see Zadeh, 1965). As a result of relaxing the underlying semantics, Degen et al.'s reconceptualization of informativity is probabilistic in nature, explaining a number of diverse phenomena in the psycholinguistics literature by adjusting several parameters in their model.

Another key feature of Degen et al.'s model (which is shared with other Rational Speech Act (RSA) models; see Frank and Goodman, 2012; Goodman and Frank, 2016) is that speakers try to be informative with regards to an internal model of the listener. While this computational architecture might suggest that over-specification is tailored to the needs of the listener, Degen and colleagues point out that such interpretation of their model is misleading: as a computational-level theory (Marr, 1982), their RSA model does not assume that speakers actively consult an internal model of a listener every time they produce an over-specific description, being agnostic as to whether human language production is based on an algorithm that explicitly involves a listener component. Thus, Degen et al. (2020) acknowledge that redundant adjectives may facilitate the listener's visual search for a referent, but they concede that “such benefits might simply be a happy coincidence and speakers might not, in fact, be deliberately designing their utterances for their addressees” (p. 617).

Like Degen et al.'s (2020) model, the visual efficiency hypothesis also moves away from the traditional pragmatic notion of informativity. However, rather than seeing over-informativity as the result of probabilistic semantics, the visual efficiency hypothesis assumes that reference is a collaborative process where both speakers and listeners try to maximize the efficiency of their interaction (Clark and Marshall, 1981; Clark and Wilkes-Gibbs, 1986). In this view, the expected difficulty of the listener's visual search for a physical referent conditions the speaker's choice of referential expression. Recent studies have offered support to this collaborative view of reference, highlighting the importance of the listener's visual search for color over-specification. Thus, participants in Rubio-Fernandez (2019; Experiments 1a and 1b) used redundant color adjectives when their listener had to find a target in a display, but not when the listener knew which object was the target in the same display. Analogously, warning participants of the risk of miscommunication resulted in an increase in color over-specification relative to a condition without such warning (Rubio-Fernandez, 2016; Experiment 2). In line with these results, van Gompel et al. (2019) have recently proposed another probabilistic model of reference generation that is sensitive to interactive factors in speaker-listener coordination, making different predictions for different referential communication tasks on the basis of an *over-specification eagerness parameter*.

The main innovation of the visual efficiency hypothesis is bringing perception and social cognition to the forefront of interactive referential communication. The role of perceptual factors in the production of referential expressions has been extensively investigated in computational models of Natural Language Generation (for reviews, see Krahmer and Van Deemter, 2012; van Deemter et al., 2012; Gatt et al., 2014). However, the architecture underlying these computational models does not include a Theory of Mind component that mimics speakers' social cognition. According to the visual efficiency

hypothesis, speakers produce referential expressions designed to facilitate the listener’s visual search for the referent. Under this formulation of efficiency, estimating the cost of producing a referential expression is not entirely *egocentric* (i.e., determined by speaker production costs), but is in fact partly *allocentric* (i.e., aimed to minimize listener costs). That means that amongst two equally informative descriptions (e.g., “The blue cup” vs. “The plastic cup”), the more efficient one would lead to faster identification of the referent in the visual context.

We have recently put forward a computational model of this account, the *Incremental Collaborative Efficiency* (ICE) model, which generates referential expressions by considering the listener’s expected visual search in real-time (Jara-Ettinger and Rubio-Fernandez, 2021, under review). To achieve this, we implemented a model of the speaker that simulates how a listener would search for an object as they process words incrementally, relying on the assumption that people can detect color from their visual periphery, but they must fixate on an object to identify its material or kind. To evaluate the ICE model’s capacity to explain reference production, we tested whether it could reproduce known qualitative patterns of over-specification: (i) speakers are more likely to over-specify color in denser visual displays (e.g., Paraboni et al., 2007; Rubio-Fernandez, 2019); (ii) this propensity, however, decreases as a function of the number of objects of the same color as the target (e.g., Koolen et al., 2013; Rubio-Fernandez, 2016); and (iii) in identical visual displays, speakers of languages with prenominal modification (such as English) are more likely to use redundant color adjectives than speakers of languages with postnominal modification (such as Spanish; Rubio-Fernandez et al., 2020; Wu and Gibson, 2021). We observed that, like human participants in referential communication studies, the ICE model’s preference for redundant color words (i) increases as a function of the number of objects in the scene, (ii) decreases with increasing monochromaticity, and (iii) is greater for prenominal adjectives than for postnominal adjectives. Critically, the ICE model predicts all these production patterns in a quantitative manner without having to add extra parameters or fit the model parameters to data.

The psychophysics of referential communication

Researchers in psychophysics investigate perceptual processes by systematically varying the properties of a stimulus along one or more physical dimensions and measuring the effect of this variation on a subject’s behavior. Computational models of color vision, for example, incorporate insights from color physics and human biology in order to predict key phenomena in color appearance (e.g., color matching, the context-dependence of color perception, or the constancy of color perception under changing lighting conditions; D’Zmura and Lennie, 1986; Maloney and Wandell, 1986; Brainard and Freeman, 1997; Irwin et al., 2000).

As illustrated in the top row of Fig. 2, psychophysics investigates the causal chain connecting a physical stimulus with an

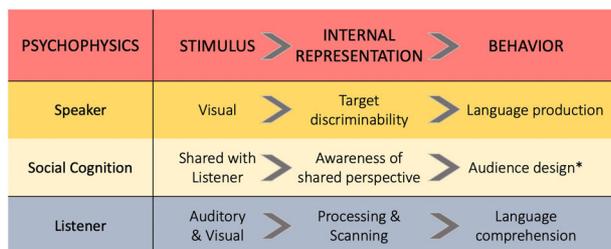
agent’s sensory experience of this stimulus, and this agent’s mental representation with their behavioral response (Lu and Doshier, 2013). Here, I want to propose a parallel analysis of referential communication that allows integrating visual perception and social cognition into a pragmatic account. Such an analysis should in turn provide us with a better understanding of the collaborative nature of reference by deconstructing referential communication from the speaker and listener perspectives (Clark and Marshall, 1981; Clark and Wilkes-Gibbs, 1986). As a starting point, a psychophysical analysis of referential communication should be applicable to the standard setup of experimental tasks: a researcher and a participant, or two participants collaborate in an interactive task in which the speaker needs to refer to a target in a display of objects, such that the listener can identify it and act on it (e.g., pick up the target and place it in a different location, or simply click on the target on a computer screen). Thus, the speaker is presented with a visual array, in which they must identify a target in order to describe it verbally, while in receiving this description, the listener is simultaneously presented with an auditory stimulus and a visual stimulus, which they must process and scan, respectively, in order to identify the intended visual target.

A major focus in psychophysics is the measurement of the *sensitivity of a perceptual system*, understood as the minimum physical stimulus (or the minimum difference between physical stimuli) that is detectable by the system (Graham, 2001; Lu and Doshier, 2013). Relatedly, successful referential communication rests on *discriminability*: in order to ensure that the listener identifies the intended referent, a speaker needs to first calculate its discriminability in the visual context; that is, they need to identify at least one property of the target that would allow the listener to distinguish it from the other objects (e.g., what kind of object it is, or what material it is made of). Calculating the discriminability of a referent in a visual context admittedly requires lower sensitivity than the perceptual tests conducted in traditional psychophysics work. Nonetheless, discriminability is fundamental to successful referential communication, more so perhaps than the long-debated notion of informativity. In order to determine how much information a speaker needs to provide their listener in order to secure unique reference, they first need to determine the discriminability of the intended referent in the visual context. Calculating a physical referent’s discriminability is therefore prior to considerations of informativity and referential choice.

The schematic analysis provided in Fig. 2 reveals another basic tenet of successful referential communication: normally, the speaker calculates the discriminability of a referent on the assumption that the listener shares their perceptual sensitivity. That is, if the speaker in a standard referential communication task identifies some unique, salient property of the referent, they can normally assume that the listener will also be able to perceive this feature. Note that this underlying assumption is not a given in every communicative situation: when communicating with a blind interlocutor, for example, a sighted speaker would need to identify tactile properties of a referent in order to unambiguously describe it (Landau and Gleitman, 2009). Relatedly, the underlying assumption that our interlocutors normally share our perceptual sensitivity touches on a long-standing debate in psycholinguistics (see Keysar et al., 2000, 2003; Tanenhaus et al., 2004): does referential communication require perspective taking?

Integrating social cognition into the psychophysics of referential communication

Speakers are supposed to engage in *audience design* when they tailor their utterances to the needs of their listeners. The key



*Informed by the Speaker’s prior sensory experience as Listener.

Fig. 2 The psychophysics of referential communication in a standard laboratory task.

question for models of audience design is whether speakers consult a discourse model of their listener in order to produce referential expressions that are comprehensible, or whether reference production results from “speaker-internal processes” (see Dell and Brown, 1991; Pickering and Garrod, 2004). The redundant use of adjectives has often been explained as a speaker-internal process, whereby a speaker encodes visually salient information without considering the listener’s perspective (for reviews, see Arnold, 2008; Davies and Arnold, 2019). For example, rather than first establishing whether there are any other cups on the table, a speaker may ask the listener for “the blue cup” simply because the cup’s color is visually salient in the context (e.g., Pechmann, 1989; Belke and Meyer, 2002; Belke, 2006; Koolen et al., 2013; Fukumura and Carminati, 2021). The received view is that, in such situations, the speaker is failing to engage in audience design because they are not tailoring their referential expression to the needs of the listener, and are relying instead on visual properties that are salient from their own perspective (for an account of egocentrism in reference comprehension, see also Keysar et al., 2000, 2003).

Here, I propose a different analysis of perspective taking in referential communication, based on the psychophysical analysis of standard tasks. In those tasks, speaker and listener normally share a physical environment—a situation also known as *co-presence* (Clark and Marshall, 1981). This is illustrated in Fig. 2, where the visual stimulus that the speaker is presented with is also shared with the listener. In this situation, the speaker is aware that they share the physical environment with their listener, and this awareness informs their audience design. Thus, whereas the perception of visual salience is indeed a speaker-internal process (e.g., registering that the cup is the only blue object on the table would not require social reasoning), a speaker’s social cognition would be recruited in mentally representing the visual environment as shared with their listener (see Graziano and Kastner, 2011; Graziano, 2013).

Central to this analysis is the notion of *mutual salience* (Rubio-Fernandez, 2019): in situations of co-presence, any property that may seem visually salient to the speaker should also be visually salient for the listener, being therefore efficient for referential communication. Thus, if a speaker perceives that the color, size or material of a target object are salient in a co-present environment, they can normally assume that these properties will also be salient for their interlocutor, without having to actively adopt their visual perspective. Relying on mutual salience is therefore easier than adopting the interlocutor’s perspective in those situations where speaker and listener have different views (see, e.g., Wardlow-Lane and Ferreira, 2008; Long et al., 2018).

As in the case of speaker and listener sharing their perceptual sensitivity, the sharedness of the physical environment is not a given in every communicative situation. When we are talking on the phone, for example, we need to describe any visually salient stimulus to our interlocutor on the other side of the line (e.g., “I’ve just seen the Empire State building turn bright pink!”). However, in situations of co-presence, it is much easier to direct our interlocutor’s attention to a visually salient stimulus; so much so that we may simply use a pointing gesture or a demonstrative pronoun (e.g., “Look at *that!*”). Importantly, distinguishing those situations where interlocutors may rely on mutual salience (e.g., cooking a meal together) from those situations that require perspective taking (e.g., giving someone directions over the phone) requires social cognition. This view is supported by anecdotal evidence with young children, suggesting that they often assume co-presence when talking on the phone, in line with studies on the development of visual perspective taking in toddlers (Moll and Tomasello, 2006; Moll et al., 2007).

As standard reference production tasks rely on co-presence and real-time interaction, while standard paradigms in social cognition research require that participants mentally simulate another person’s visual perspective, or imagine someone else’s actions (Ruby and Decety, 2001; Bio et al., 2018), it could be argued that referential communication need not recruit social cognition. However, Deschrijver and Palmer (2020) have recently proposed that *mental conflict monitoring* is one of the main functions of social cognition—perhaps even more fundamental than the ability to attribute mental states to others. Through mental conflict monitoring, we keep track of how others think differently from us, which is a key skill to successfully navigate the social world. According to Deschrijver and Palmer, when two people share the same perspective (which in Theory of Mind research is normally set up as a true-belief scenario, where both parties are right about a certain state of affairs), they still engage in mental conflict monitoring. That is, humans estimate the degree to which they share their own perspective with others, even in situations where their perspectives are aligned. This account of social cognition supports the psychophysical analysis of referential communication tasks schematized in Fig. 2, since establishing mutual salience in situations of co-presence would be the output of mental conflict monitoring.

Efficient audience design

Ferreria (2019) has recently proposed a mechanistic framework for audience design where speakers may be sensitive to the discourse needs of their interlocutors without having to engage in an effortful process of perspective taking (for earlier precursors of this idea, see Jaeger and Ferreira, 2013; Kurumada and Jaeger, 2015; Buz et al., 2016). This new model challenges traditional assumptions about pragmatics, which is often viewed as a cognitively taxing component of language production and language comprehension (for discussion, see Geurts and Rubio-Fernandez, 2015). The crucial innovation in Ferreira’s framework is that feedforward language production allows speakers to implement audience design strategies in a relatively automatic way (i.e., without the involvement of executive control). These tacit strategies are what Ferreira calls *feedforward audience design*. Importantly, this form of audience design is heavily limited since tacit strategies, learned through communicative experience, must be available prior to grammatical encoding in order to drive feedforward audience design. However, given time and resources, speakers are also able to evaluate the potential communicative effect that a particular utterance will have on a given listener—which in Ferreira’s framework is termed *recurrent processing audience design*.

In line with this work, I have recently argued that speakers rely on perceptual contrast as a visual heuristic in the production of referential expressions (Rubio-Fernandez, 2019; Long et al., under review), and that this visual heuristic is a learned strategy for feedforward audience design (Ferreria, 2019). Thus, when speakers detect a sufficient degree of perceptual contrast in a visual context (e.g., when objects clearly vary in color, size or shape), they would use the corresponding adjective in their referential expression, even if it may not be necessary to establish unique reference. This is an efficient referential strategy because it allows speakers to use perceptual contrast as a proxy for the calculation of discriminability. Thus, relying on perceptual contrast reduces the cost of producing an appropriate referential expression (e.g., speakers do not need to scan the display for competitors; see Davies and Kreysa, 2017) and facilitates the listener’s visual search for the referent by providing contrastive information (Long et al., under review).

In addition, in situations of co-presence, a speaker may engage in audience design by relying on their *prior experience as listeners* (for empirical evidence, see Vogels et al., 2020). Consider, for example, the difference between color and material adjectives (e.g., “the blue cup” vs. “the plastic cup”). From a statistical learning perspective, speakers should produce more color adjectives than material adjectives simply because color words are more frequent in their language input (Christiansen, 2019). However, when listeners process color and material adjectives, they are not only exposed to certain word frequencies and statistical regularities: they are also experiencing, first hand, the differential perceptual effort that is required to discriminate an object by its color or by its material.

In most circumstances, it is easier to identify an object by its color than by its material (for eye-tracking evidence and computational modeling, see Jara-Ettinger and Rubio-Fernandez, 2021, under review). Therefore, from a psychophysics perspective, speakers’ tendency to over-specify color (rather than material) would result, to some extent, from their prior sensory experience as listeners. In this view, speakers’ reliance on their prior communicative experience could be understood as a form of perspective taking that informs their feedforward audience design (Ferreria, 2019).³

Previous work on the visual efficiency of over-specification tried to establish whether speakers are sensitive to their listener’s needs (see Rubio-Fernandez, 2016, 2019; Rubio-Fernandez et al., 2020). Building on the results of those studies, here I will assume that speakers are cooperative, and will employ language comprehension and language production tasks to investigate whether redundant color adjectives are efficient for referential communication.

The effect of redundant modification on visual search

A number of psycholinguistic studies since the 80’s have investigated the effect of redundancy on reference resolution, with mixed results: some studies have shown that redundancy hinders referential communication, while others revealed that adding extra information facilitates the identification of the intended referent. Supporting the traditional view that redundancy is inefficient, several psycholinguistic studies have shown that redundant descriptions can impair communication (Engelhardt et al., 2006, 2011; Davies and Katsos, 2013). However, upon closer inspection of their materials, the redundant information used in these studies may not have had discriminatory value that could facilitate visual search. For example, Engelhardt et al. (2006) used complex descriptions such as “Put the apple on the towel in the box”, where the over-informative phrase “on the towel” could refer to the current location of the apple or to its destination, creating a temporary ambiguity that should delay visual search rather than facilitate it.

More directly comparable to the present study, Engelhardt et al. (2011) reported that redundant descriptions like “Look at the red square” delayed target identification in displays with two shapes, even though color was a distinctive property of the target. Engelhardt and colleagues concluded that redundant information impairs communication, yet their results could be related to the sparse displays they employed. This is relevant to the visual efficiency hypothesis because a property’s discriminatory value does not solely depend on its distinctiveness, but also on the density of the display, with redundant color words being more efficient in denser displays where the listener’s visual search is harder (Rubio-Fernandez et al., 2020). Supporting this view, a number of studies have shown that redundant referential expressions are more frequent in denser displays (e.g., Paraboni et al., 2007; Clarke et al., 2013; Koolen et al., 2015; Gatt et al.,

2017). In particular, Rubio-Fernandez (2019) observed that in two-shape displays, people over-specified color around 20% of the time, whereas they did so over 80% of the time in displays with eight shapes. It is therefore an open question whether redundant color words would facilitate visual search in displays with more than two shapes, where redundant color words are more frequently produced.

Davies and Katsos (2013) also observed delayed response times with redundant descriptions such as “Pass me the small star” in displays of four objects. However, Davies and Katsos pretested their materials and only used properties that were not frequently over-specified in a language production task. Again, this leaves open the possibility that properties that are frequently over-specified (such as color) may have facilitated participants’ visual search.

Challenging the conclusions of the above studies and supporting the hypothesis that redundant adjectives can be visually efficient, a number of studies have shown that redundant properties such as color, size and position can sometimes reduce referent identification times (Sonnenschein and Whitehurst, 1982; Mangold and Pobel, 1988; Paraboni et al., 2007; Arts et al., 2011; Paraboni and Van Deemter, 2014). However, these studies used off-line tasks (e.g., participants had to read a target description before searching for it in a display, or they had to navigate a virtual environment following over-informative spatial descriptions). Engelhardt and Ferreira (2016) challenged the reliability of these results on the grounds that processing a description ahead of time could allow participants to build a more detailed mental representation of the target that could later facilitate its identification, yet this facilitation may not be observed in real-time language processing. In order to address this legitimate criticism, monitoring visual search during real-time language processing would offer a better test of the visual efficiency hypothesis.

Prior work on visual attention has shown that real-time language interpretation can constrain visual perception; in particular, the incremental processing of a spoken instruction results in the listener’s search being less affected by the number of distractors (as in *feature search tasks*; Treisman and Gelade, 1980; Spivey et al., 2001; Reali et al., 2006; Lupyán and Spivey, 2010). Eye-tracking studies investigating real-time language interpretation have compared the processing of modified instructions when color was necessary to distinguish two competitors and when it was redundant (Sedivy, 2003, 2004; Grodner and Sedivy, 2011; Tourtouri et al., 2019; see also Engelhardt et al., 2011). For example, in a recent eye-tracking study using the visual-world paradigm, Tourtouri et al. (2019) compared the processing of “Find the blue ball” in displays with one or two balls, and concluded in view of their results that redundancy facilitates the listener’s processing. It is important to note that by comparing contrastive and descriptive uses of adjectives, these studies investigated the effect of redundancy on adjective processing. However, the experimental design used in the above studies does not address the issue of whether redundant color words can facilitate the visual search for a referent. In order to investigate that question, listeners’ eye movements and response times should be compared during the processing of redundant and minimal instructions (e.g., “Click on the blue star” vs. “Click on the star”) in the same visual display.

Comparing the processing of a modified description in displays with or without competitors vs. comparing the processing of modified and unmodified descriptions in a visual display without competitors is an important methodological difference that holds the key to understanding the nature of referential over-specification. Thus, comparing the processing of modified descriptions when color is necessary or redundant cannot explain

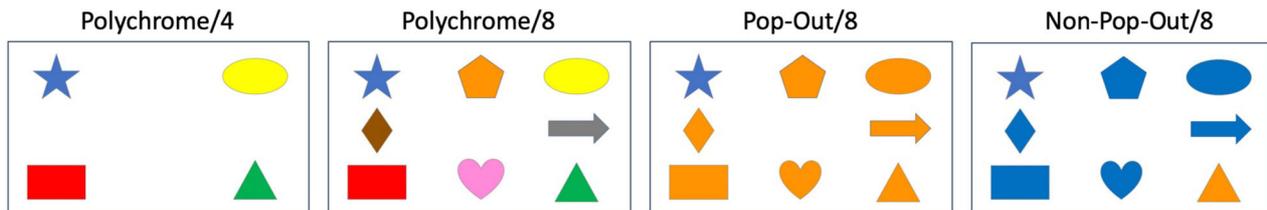


Fig. 3 Sample items from the four conditions in Experiment 1. Eye movements and response times were used to measure target identification during language processing. The target in all four displays was the star, and the instructions included bare definite descriptions (“Click on the star”) and redundant color adjectives (“Click on the blue star”).

why people use color words redundantly. When presented with a display of shapes with a single star, for example, speakers do not have the choice to “throw in” a second star and make their color adjective contrastive. However, that is essentially what previous eye-tracking studies compared in order to explain color over-specification (Sedivy, 2003, 2004; Grodner and Sedivy, 2011; Tourtouri et al., 2019; see also Engelhardt et al., 2011). The choice that speakers face when presented with a display of shapes with a single star is whether or not to *mention its color*, and the question therefore remains: would asking for “the blue star” facilitate the listener’s visual search relative to the shorter description “the star”?

To summarize, a considerable number of studies have investigated the effect of redundancy on reference resolution with mixed results. Whereas most studies have shown a facilitatory effect, the use of off-line tasks that do not tap real-time language processing, or inadequate eye-tracking materials that do not map onto a speaker’s choice when referring to a singleton object make it impossible to derive any firm conclusions about the effect of redundancy on a listener’s visual search for a referent (for related arguments, see Degen et al., 2020). The present study tries to overcome the methodological limitations of previous studies in an eye-tracking study on the role of discriminability in reference resolution, as well as reference production.

An empirical test of the visual efficiency hypothesis

Given the fundamental role that discriminability plays in referential communication, an eye-tracking study was conducted to investigate whether redundant color adjectives facilitate target identification in contexts where color is distinctive of the target. Eye movements and response times were compared when the instructions included a redundant color adjective (e.g., “Click on the blue star”) as opposed to a bare noun (e.g., “Click on the star”). Based on Rubio-Fernandez (2019; Experiments 1a and 1b), four display types were used as visual materials: polychrome displays of four geometrical shapes (Poly/4) or eight geometrical shapes (Poly/8), and pop-out displays of eight geometrical shapes where the target was either the only shape of its color (Pop-Out/8), or the same color as most shapes (Non-Pop-Out/8; see Fig. 3 for sample items). Rubio-Fernandez (2019) observed that speakers used redundant color adjectives in conditions analogous to Poly/4, Poly/8 and Pop-Out/8, but produced zero rates of redundant color adjectives in monochrome displays where color was not distinctive of the target. These results supported the view that redundant speakers are cooperative, producing color adjectives when color is distinctive and may facilitate the listener’s search for the target. The question remains, however, as to whether redundant color adjectives do indeed facilitate target identification for the listener. This was the aim of Experiment 1, which tried to directly test the visual efficiency hypothesis using eye-tracking during real-time language processing.

The main prediction that was tested in Experiment 1 is the reverse effect of redundant color adjectives on visual search

depending on color discriminability (i.e., facilitating vs. hindering visual search). Contrary to the standard pragmatics view of redundancy as inefficient, the visual efficiency hypothesis predicts that redundant color adjectives will only hinder visual search when color does not discriminate the referent, whereas it should facilitate target identification for the listener when it is distinctive. It follows from this prediction that participants should identify “the blue star” faster than “the star” in Poly/4, Poly/8, and Pop-Out/8, whereas the reverse should hold for the Non-Pop-Out/8 condition. Crucially, “blue” is redundant in all four display types, so according to the standard view, all these uses are pragmatically infelicitous.

Experiment 2 investigated the role of target discriminability on the production of redundant color adjectives. Rubio-Fernandez (2019) observed that people produced redundant color adjectives in polychrome displays, but not in monochrome displays, contrasting two extreme cases in target discriminability: the color of the target was either unique in the display, or uniform across the display. Such a comparison, however, does not reveal the degree to which speakers are sensitive to target discriminability, and whether this perceptual sensitivity affects their production of redundant color adjectives. This is an important question when investigating the psychophysics of referential communication. Therefore, in the second experiment, the discriminability of the target color parametrically varied in displays of 9 shapes, ranging from maximal discriminability (i.e., the target color was unique in the display) to zero discriminability (i.e., all shapes were the same color as the target; see Fig. 4 for sample items).

Experiment 1

Methods

Participants. Ethics approval for the study was obtained from the UCL REC Committee. All participants gave signed informed consent prior to the start of the task. A group of 40 students from University College London took part in the first experiment. They were native speakers of English and participated for monetary compensation. All participants reported having normal color vision.

Design and procedure. The materials included 8 displays from each of the following conditions: Poly/4, Poly/8, Pop-Out/8, and Non-Pop-Out/8 (for sample displays, see Fig. 3). Critical trials were presented twice during the study, for a total of 64 critical trials. An extra eight slides (two from each condition) were used as practice trials. In addition, another 64 slides, similar to the critical ones but repeating non-target colors and shapes, were used as filler trials in order to avoid that participants would develop strategies. Practice and filler trials were not analyzed. There were 136 trials in total. Color was redundant in all trials.

The target colors were: blue, brown, green, gray, orange, pink, purple, red and yellow; and the target shapes were: circle, cross, diamond, heart, pentagon, rectangle, square, star and triangle. All shapes were different in each display (making the use of color

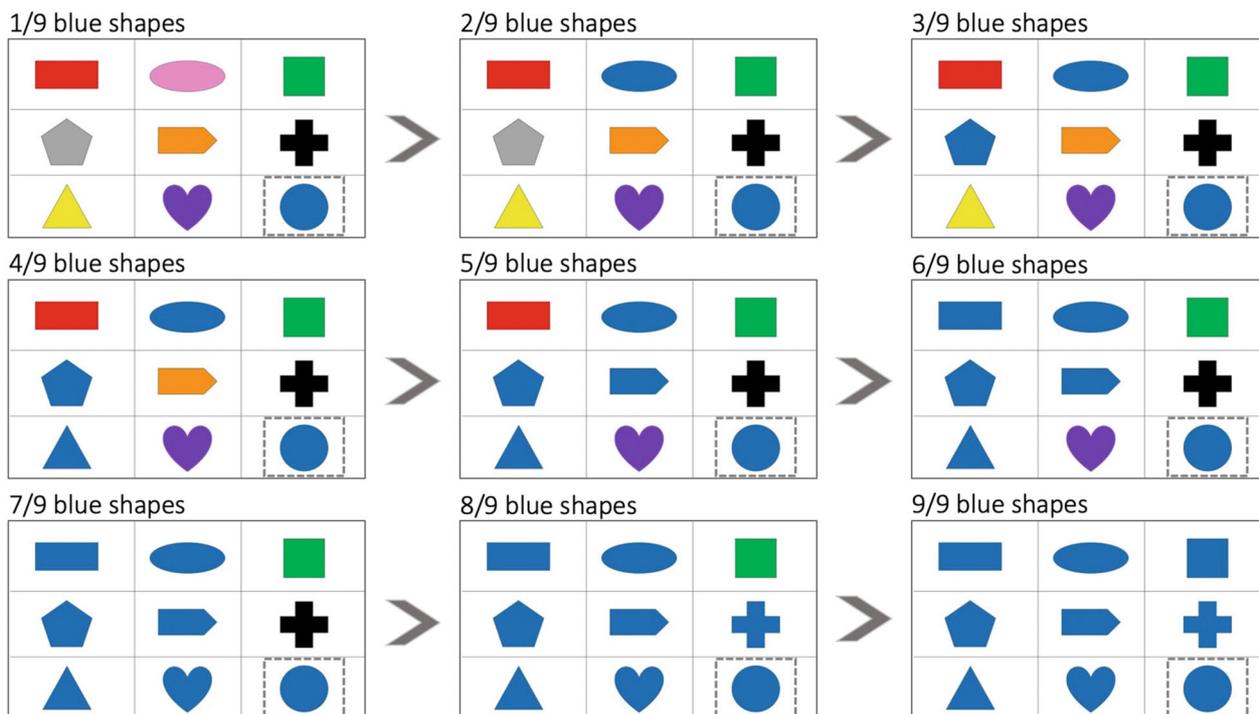


Fig. 4 Sample items from the nine conditions used in Experiment 2 (1-9 shapes in the target color). Participants had to refer to the target shape that appeared inside a box. Rates of redundant color adjectives (e.g., “the blue circle”) were calculated as a function of color discriminability.

redundant). Color was distinctive of the target shape in three of the critical conditions (Poly/4, Poly/8 and Pop-Out/8), whereas it did not distinguish the target shape in the fourth condition (Non-Pop-Out/8). The aim of the Non-Pop-Out/8 condition was twofold: first, it prevented participants from anticipating the target in the Pop-Out/8 condition (since otherwise participants could learn that the target would be the singleton shape in all pop-out displays). Second, it allowed comparing the effect of redundant color adjectives when the target color was distinctive and when it was not, serving as a baseline for the Poly/4, Poly/8, and Pop-Out/8 conditions. Thus, Experiment 1 investigated the real-time effect of redundant color adjectives relative to shorter descriptions as a function of color discriminability, serving as a direct test of the visual efficiency hypothesis.

The displays were on the screen for 400 ms before the instruction started. The instructions were recorded by a male speaker of British English, who did not stress the adjectives contrastively. Each display was paired with two different instructions: one redundant, including the color of the target (e.g., “Click on the blue star”) and a minimal one, only mentioning the shape of the target (e.g., “Click on the star”). Phonological competitors of the target were avoided both for the color and shape words, as they could create interference during processing. The visual materials were presented twice during the task, once paired with a redundant instruction and once with a minimal instruction. To minimize recency effects, trials were split into two blocks, one in the Redundant condition and another one in the Minimal condition. Between each block there was a 20 s break to give participants an opportunity to rest. The order of the items was randomized within each block and the order of the blocks was counterbalanced in two separate lists of materials, which were randomly and evenly distributed amongst participants.

Participants were told that they were going to listen to a series of instructions to click on different shapes and their task was to click on the target as fast and accurately as possible. In between

trials, they had to click on a cross in the center of the screen, which ensured that the mouse cursor was always in the center of the screen at the start of each trial.

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. The contact-free set up of the system allowed free head-movement during eye tracking. The experiment lasted ~15 min.

Processing measures of (i) delay of first target fixation, (ii) percentage of target fixations in the critical time window, and (iii) response times (RTs) were collected. RTs were measured from the onset of the instruction, and first target fixations were measured from the onset of the first word in the instruction that could distinguish the target (i.e., the color adjective in the Redundant condition and the shape noun in the Minimal condition). In the Redundant condition, the critical time window was calculated from the onset of the adjective until the offset of the noun. In the Minimal condition, the critical time window started at noun onset and included an equivalent silence after the noun (i.e., [ADJECTIVE + NOUN] = [NOUN + SILENCE]; $M = 900$ ms). The critical time windows in the Redundant and Minimal conditions were matched in length individually for each item.

It was predicted that, if redundant color adjectives are visually efficient, an advantage should be observed in the Redundant condition relative to the Minimal condition when color is a distinctive property of the target (Poly/4, Poly/8, and Pop-Out/8). However, when color is not distinctive (Non-Pop-Out/8), a disadvantage should be observed in the Redundant condition because the adjective creates a temporary ambiguity during processing. For example, when processing the color adjective in the instruction “Click on the blue star”, participants could identify the target in conditions Poly/4, Poly/8, and Pop-Out/8 because the target color is distinctive (see Fig. 3 for sample items). However, when processing the same instruction in the Non-Pop-

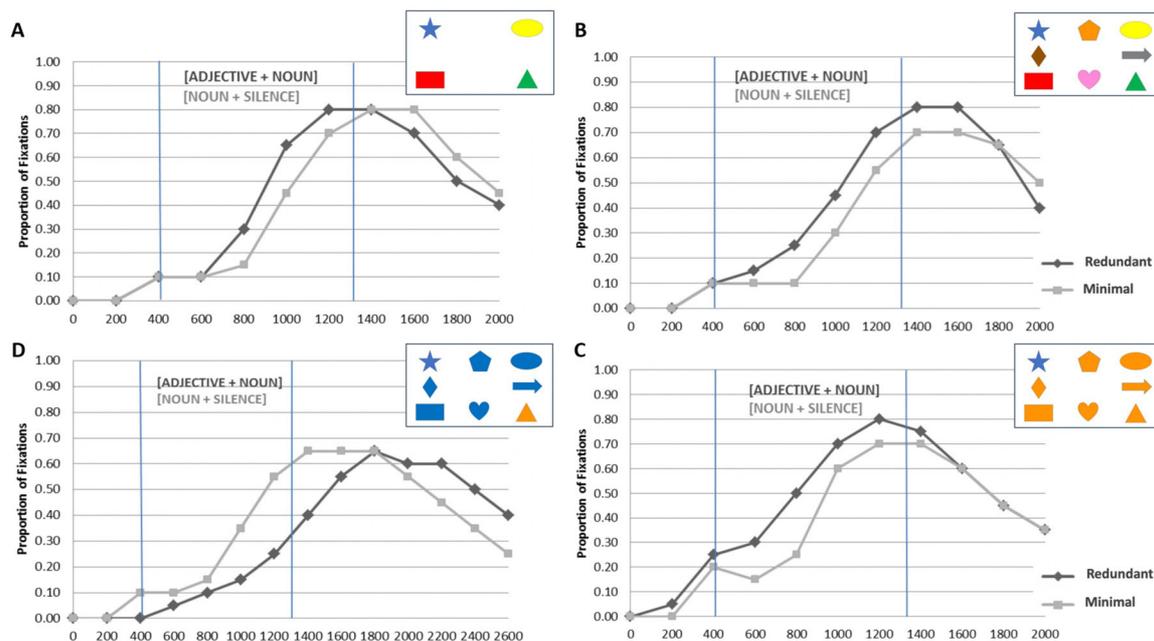


Fig. 5 Proportion of target fixations across time in the four visual conditions of Experiment 1. A–D Proportions of target fixations during the processing of redundant instructions (“Click on the blue star”) and minimal instructions (“Click on the star”) in conditions (A) Polychrome/4, (B) Polychrome/8, (C) Pop-Out/8, and (D) Non-Pop-Out/8. The target in all four displays was the blue star, and color was distinctive of the target in (A), (B), and (C), but was non-distinctive of the target in D. Redundant color adjectives were therefore predicted to have a facilitatory effect in (A), (B), and (C) relative to the minimal instructions, whereas it was predicted that color should delay target identification in D because it creates a temporary ambiguity. Blue vertical lines mark the critical time window (i.e., [ADJECTIVE + NOUN] in the Redundant condition and [NOUN + SILENCE] in the Minimal condition).

Out/8 condition, participants would not be able to identify the target when hearing “blue”, having to wait for the shape noun to disambiguate the instruction. The visual efficiency hypothesis therefore makes reverse predictions depending on a target’s color discriminability. However, since color was redundant in all displays, standard pragmatic accounts based on informativity do not make different predictions depending on color discriminability.

Results. Delayed eye movement (happening only once the participant had clicked on the shape) and erroneous responses led to the removal of 38 first target fixations (3% of data) and 25 RTs (2% of data). Proportions of target fixations during the processing of redundant and minimal instructions in the four display types are plotted in Fig. 5A–D.

Linear mixed effects regression was used for each of the processing measures, modeling the fixed effects of Display and Instruction plus their interaction, including Participants and Items as random effects, and the maximal random effect structure (for model outputs, see Table 1). Both predictors were dummy coded with the Non-Pop-Out/8—Minimal condition as the reference level because this was the only condition where color was not a distinctive property of the target. It was therefore predicted that color adjectives would have the reverse effect in the other display types (Poly/4, Poly/8, and Pop-Out/8; see Fig. 6A–C). 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).

For each of the three processing measures, the model revealed a main effect of Instruction in Non-Pop-Out/8: redundant instructions resulted in slower first target fixations ($B = 250.18$; $t = 5.35$), lower percentages of target fixations during the critical time window ($B = -17.92$; $t = -4.29$) and slower response times ($B = 341.75$; $t = 5.22$) than the minimal instructions. The Display \times Instruction interactions were as predicted: relative to the

displays where color was not distinctive, redundant color adjectives sped up first target fixation in Poly/4 ($B = -353.17$; $t = -7.39$), Poly/8 ($B = -400.11$; $t = -6.92$) and Pop-Out/8 ($B = -380.90$; $t = -7.58$). Redundant color adjectives also increased the percentage of target fixations during the critical time window in Poly/4 ($B = 24.54$; $t = 5.02$), Poly/8 ($B = 34.34$; $t = 5.79$), and Pop-Out/8 ($B = 28.77$; $t = 6.62$). Finally, redundant color adjectives reduced response times when color was distinctive of the target (Poly/4: $B = -421.41$; $t = -4.60$; Poly/8: $B = -468.34$; $t = -4.55$; Pop-Out/8: $B = -488.86$; $t = -6.43$).

According to the visual efficiency hypothesis, redundant instructions should have a facilitatory effect on visual search relative to minimal instructions in Poly/4, Poly/8 and Pop-Out/8. In order to determine if there was a significant main effect of Instruction in each of these conditions, the same linear mixed effects regression model (including fixed effects of Display and Instruction plus their interaction, Participants and Items as random effects and the maximal random effect structure) was run again on each of the three measures (First Target Fixation, Percentage of Target Fixations and RTs), but changing the reference level to the other display types. That is, rather than using the Non-Pop-Out/8—Minimal condition as the reference level, the predictors were dummy coded with Poly/4—Minimal, Poly/8—Minimal and Pop-Out/8—Minimal as reference level, in three additional runs of the same model on each data set.

The main effect of Instruction was significant on first target fixations, revealing faster target identification in the Redundant condition than in the Minimal condition in Poly/4 ($B = -103$; $t = -2.51$), Poly/8 ($B = -149.45$; $t = -2.39$) and Pop-Out/8 ($B = -129.81$; $t = -2.62$). Instruction also had a significant effect on the percentage of target fixations in the critical time window, revealing higher rates of target fixations in the Redundant condition than in the Minimal condition in Poly/8 ($B = 16.42$; $t = 3.71$) and Pop-Out/8 ($B = 10.85$; $t = 4.03$), and a marginally significant difference in Poly/4 ($B = 6.63$; $t = 1.98$). As for RTs,

Table 1 Results of linear mixed effect models for each processing measure, positing fixed effects of Display and Instruction plus their interaction, with random effects of Subject and Item, and the maximal random effect structure.

Measure	Fixed effects	Estimate	Std. error	t-value
First target fixation	Intercept	1252.78	96.67	12.96
	Poly/4	-165.46	106.60	-1.55
	Poly/8	-65.36	119.51	-0.55
	Pop-Out/8	-238.31	111.40	-2.14
	Redundant	250.18	46.79	5.35
	Poly/4 vs. Redundant	-353.17	47.83	-7.39
	Poly/8 vs. Redundant	-400.11	57.83	-6.92
	Pop-Out/8 vs. Redundant	-380.90	50.24	-7.58
Percentage of target fixations	Intercept	39.17	6.62	5.92
	Poly/4	15.01	7.66	1.96
	Poly/8	2.24	8.46	0.27
	Pop-Out/8	15.83	6.81	2.33
	Redundant	-17.92	4.18	-4.29
	Poly/4 vs. Redundant	24.54	4.89	5.02
	Poly/8 vs. Redundant	34.34	5.93	5.79
	Pop-Out/8 vs. Redundant	28.77	4.34	6.62
Response times	Intercept	2042.51	102.85	19.86
	Poly/4	-188.12	128.92	-1.46
	Poly/8	-53.39	128.17	-0.42
	Pop-Out/8	-231.58	129.00	-1.80
	Redundant	341.75	65.47	5.22
	Poly/4 vs. Redundant	-421.41	91.67	-4.60
	Poly/8 vs. Redundant	-468.34	103.03	-4.55
	Pop-Out/8 vs. Redundant	-488.86	75.98	-6.43

Non-Pop-Out/8—Minimal was dummy coded as the reference level because Non-Pop-Out/8 was the only display type where target color was not distinctive, so reverse results were predicted in the other three conditions.

participants were faster to click on the target in the Redundant are a more coarse-grained processing measure, it is not

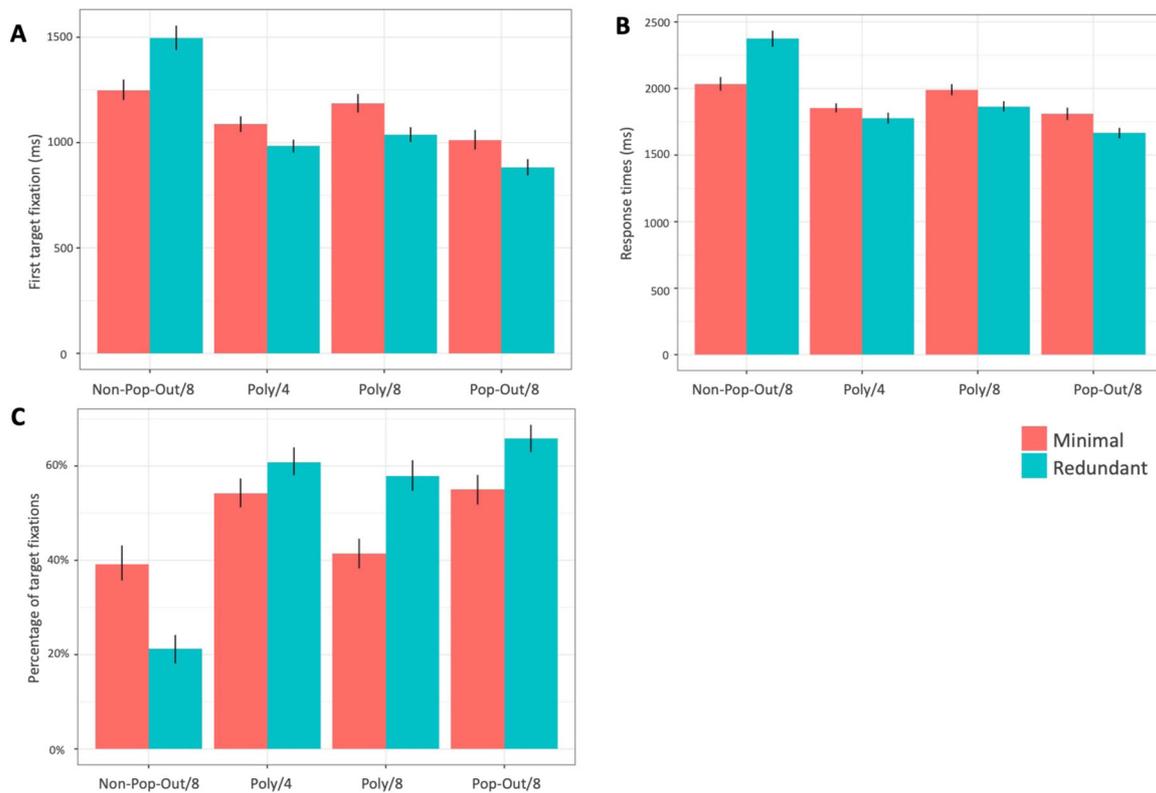


Fig. 6 Descriptive statistics from Experiment 1. A Mean delay of first target fixation and **B** mean RTs by Display and Instruction type. **C** Mean percentages of target fixation during the critical time window. Non-Pop-Out/8 was the only condition where color did not discriminate the target. Error bars indicate 95% bootstrapped confidence intervals.

condition than in the Minimal condition, but the effect of Instruction was marginally significant in Poly/8 ($B = -126.47$; $t = -1.84$) and Pop-Out/8 ($B = -147.11$; $t = -2.24$), and did not reach significance in Poly/4 ($B = -79.66$; $t = -1.25$). Since RTs

surprisingly that stronger results were observed with eye-tracking measures.

The results of Experiment 1 showed, for the first time, that redundant color modification facilitates visual search for the

referent relative to shorter descriptions. The key role of color discriminability was also confirmed, with redundant color adjectives facilitating or delaying visual search depending on the discriminability of the color in the visual display. Experiment 2 further investigated the degree to which speakers are sensitive to color discriminability when producing redundant modification.

Experiment 2
Methods

Participants. A total of 396 participants were recruited through Prolific. Geographical area was set to the UK, age was limited to 18 to 40 years, and only monolingual speakers of English were allowed to participate. These demographics were selected to match those in Experiment 1. Unlike the first experiment, Experiment 2 had to be run online, rather than in the lab, because of the COVID pandemic.

Design and procedure. Materials consisted of 36 displays of 9 different shapes. The shapes were the following: arrow, circle, cross, heart, oval, pentagon, rectangle, square, and triangle (see Fig. 4 for sample displays). The targets were a blue circle, a green square, a red rectangle and a yellow triangle, each featuring in nine displays. Targets were always placed in a corner and marked by a dashed frame. Displays systematically varied in the discriminability of the target color, which ranged from one shape (unique target-color, maximal discriminability) to nine shapes (monochrome display, zero discriminability). The distribution of the shapes in the same color as the target (henceforth, *target-color shapes*) was different for each item.

As the task was administered online, participants had time to develop strategies during the task (e.g., if they noticed that monochromaticity was manipulated across trials). To avoid such strategies, participants performed a single-trial task (with 11 participants being presented with each display in the materials, for a total of 36 displays and 396 participants). Participants had to ask an imaginary interlocutor to click on the target shape, for which they had to complete the instruction “Click on the _____”. Participants were asked to express themselves naturally, avoiding definitions (e.g., “The figure with four equal straight sides and four right angles” instead of “the square”) or spatial coordinates (e.g., “The top right corner”), which would prevent the use of color adjectives. Participants were told that their virtual partner did not know which shape was the target because in the partner’s display, the target did not appear inside a box. This was done in order to ensure that participants would produce sufficiently informative descriptions.

Results. Of the 396 responses collected, 22 had to be discarded (5.6%) because they included underinformative instructions (e.g., “Click on the lovely shape”) or referred to the location of the target, disregarding the task instructions (e.g., “Click on the top left cell”). The remaining 374 responses were coded for the presence or absence of a redundant color adjective, corresponding with redundant and minimal instructions, respectively (e.g., “The blue circle” vs. “The circle”). Only responses including a color adjective plus a shape noun were coded as redundant (i.e., “The blue one” was coded as minimal)⁴. The percentage of redundant and minimal descriptions are plotted in Fig. 7 as a function of target-color discriminability (i.e., 1–9 target-color shapes, from maximal to zero discriminability).

Using logistic mixed effects regression, the binary outcome variable of presence/ absence of Modification (1 = redundant color adjective, 0 = bare shape noun) was modeled with Condition (number of target-color shapes) as a continuous predictor. The maximal random effect structure was used, with

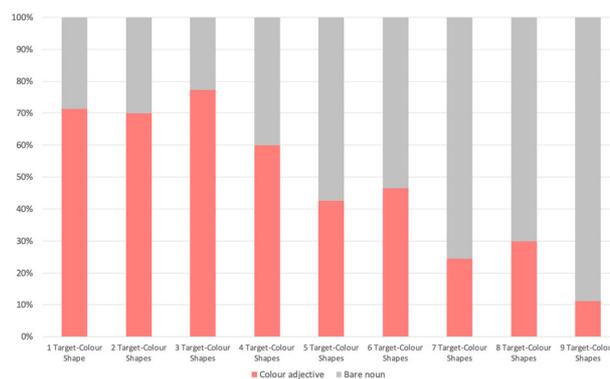


Fig. 7 Percentage of redundant color adjectives and bare nouns as a function of target-color discriminability. Target-color discriminability ranged from maximal (one Target-Color Shape: the target has a unique color) to zero (nine Target-Color Shapes: all nine shapes are the same color as the target).

by-item random intercepts. The 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). The analysis revealed a significant effect of Condition ($B = -0.3688$, $SE = 0.0530$, $p < .001$), with a decrease in modification as the target color became less discriminable (see Fig. 8). The results of Experiment 2 therefore confirm that speakers are highly sensitive to the discriminability of a target referent, and that this perceptual sensitivity directly informs their production of redundant color adjectives.

General discussion

A new pragmatic analysis of referential communication was proposed, which tries to integrate visual perception and social cognition by borrowing insights from psychophysics. Psychophysics investigates the relationship between physical stimuli and mental phenomena, providing a promising framework for the study of interactive referential communication. Standard referential communication tasks rely on the co-presence of speaker and listener, and their real-time interaction with objects. Thus, successful referential communication rests not only on the interlocutors’ linguistic abilities, but also on their shared perception of the environment. I have therefore argued that the speaker’s perception of the environment as shared with the listener (what is also known as co-presence; Clark and Marshall, 1981) is a form of social reasoning, in line with Deschrijver and Palmer’s (2020) mental conflict monitoring. When speaker and listener have conflicting perspectives on a visual display, engaging in audience design requires that speakers actively adopt the listener’s perspective, which is cognitively costly (e.g., Wardlow-Lane and Ferreira, 2008; Long et al., 2018). However, in situations of co-presence, speakers can engage in audience design by relying on mutual salience (i.e., the assumption that visual properties that are salient for the speaker will also be salient for the listener by virtue of being co-present; Rubio-Fernandez, 2019).

In line with Ferreria (2019) feedforward audience design, I have also argued that speakers can rely on their prior sensory experience as listeners to produce referential expressions that facilitate the listener’s visual search for the referent. For example, having experienced the value of contrastive adjectives for successful referential communication, when speakers perceive a sufficient degree of contrast in a visual display (e.g., whether the objects vary by color, size or shape), they may produce the corresponding adjective without scanning the display for competitors (Long et al., under review). In this view, contrast perception

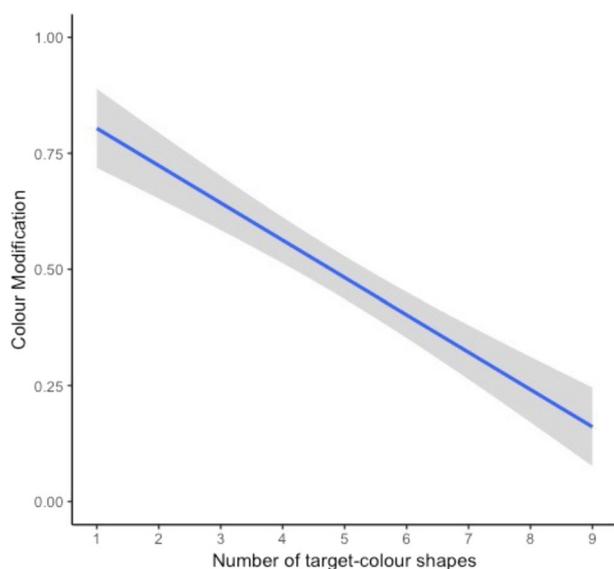


Fig. 8 Regression line showing the main effect of Condition (number of target-color shapes) on the production of redundant color adjectives. The shaded band around the regression line represents a 95% confidence region for the regression fit.

works as a proxy for the calculation of discriminability, resulting in an efficient form of feedforward audience design.

The present study confirmed the visual efficiency hypothesis by investigating the role of discriminability in reference comprehension and production in two experiments using similar visual displays. In Experiment 1, results from two eye-tracking measures and response times confirmed that redundant color adjectives can speed up target identification relative to shorter descriptions, provided color is distinctive of the target. Even if redundant color descriptions may require more effort to produce and a longer time to process, they were efficient in facilitating the listener's visual search for the referent. This finding highlights the importance of perceptual factors (and not only information quantity) when calculating a referential expression's efficiency (Rubio-Fernandez, 2016, 2019; Rubio-Fernandez et al., 2020). Here, participants identified "the blue star" faster than "the star" when the star was the only blue shape in the display, but they were faster to follow the shorter instruction when most shapes were blue. As color was redundant in both situations, standard pragmatic accounts based on information quantity fall short of explaining this pattern of results. However, a pragmatic analysis based on discriminability predicts reverse results in the two conditions.

In Experiment 2, speakers produced higher rates of redundant color adjectives the more discriminable the target color was in the display. These results confirm that participants in referential communication tasks are not only sensitive to extreme discriminability conditions (i.e., displays where the color of the target is unique vs. displays where all shapes are the same color; Rubio-Fernandez, 2016, 2019): here I showed that, on average, speakers are sensitive to the discriminability of a target color along a perceptual continuum, deploying this high sensitivity in their use of redundant color adjectives. The results of these two experiments therefore confirm that perceptual discriminability determines both speakers' use and listeners' processing of redundant color adjectives, hence preserving the efficiency of their referential communication.

Experiments 1 and 2 used similar displays of geometrical shapes to investigate the role of discriminability in reference

production and reference comprehension. Thus, speakers' sensitivity to the discriminability of a target's color in their use of redundant color adjectives is in line with the advantage that redundant color adjectives give listeners when color is distinctive. However, because of the ongoing pandemic, language production data was collected online in the written modality, rather than elicited in the lab from real-time interaction (for lab studies with similar materials, see Rubio-Fernandez, 2019; Long et al., 2020; Rubio-Fernandez et al., 2020). For a more accurate investigation of the relationship between reference production and reference comprehension in interactive communication, future studies should collect production data in the lab, and try to correlate the degree to which participants use redundant color adjectives under different discriminability conditions with the relative advantage that color gives listeners when identifying a target in the same visual displays. In addition, Experiment 2 was designed as a single-trial task (in order to avoid that participants could develop strategies), so speakers' sensitivity to discriminability was measured across participants rather than within. Future reference production studies should also collect multiple data points from each participant in order to determine whether individual differences in discriminability affect color adjective use.

Previous eye-tracking studies compared the processing of modified descriptions (e.g., "Find the blue ball") in conditions where adjectives were necessary (e.g., in a display with two balls) and conditions where adjectives were redundant (e.g., in a display with a single ball; Tourtouri et al., 2019; see also Sedivy, 2003, 2004; Grodner and Sedivy, 2011). However, the processing advantage that modification may reveal when it is redundant vs. when it is contrastive does not speak to the issue of whether redundant modification facilitates the listener's search for a referent. The answer to that question requires comparing listeners' processing of minimal and redundant descriptions in the same visual display. The present study is therefore the first to show that listeners can indeed identify a target faster following a color-redundant description (e.g., "the blue star") than a minimal description (e.g., "the star"), provided the color of the target is distinctive.

A recent eye-tracking study by Fukumura and Carminati (2021) investigated the processing of twice-modified descriptions (e.g., "Click on the spotty green bow" in a display with four different bows) and manipulated whether one or both adjectives were necessary to identify the referent. The mention of color as a second adjective delayed response times relative to a more concise pattern-only description, even though it increased fixations to the target. The authors conclude that listeners attend to all the attributes mentioned in a referring expression before selecting a referent, and this is why over-specification can hinder comprehension. While this conclusion may apply to visual displays of four objects of the same kind (in which identifying a specific object may be harder), the results of Fukumura and Carminati's study contradict the findings of previous eye-tracking studies using displays with different kinds of objects. Thus, Eberhart et al. (1995) and Rubio-Fernandez and Jara-Ettinger (2020) showed that listeners identified a target as soon as they had sufficient information to do so (e.g., when they heard a color adjective) and disregarded later information if it applied to multiple objects (e.g., if there were several objects that matched the noun).

The results of the present study also contradict Fukumura and Carminati's (2021) predictions for unmodified descriptions (which they did not test in their experiments): "if the noun alone discriminates fast in the context, additional color mention may not speed up referent selection compared to a shorter, noun-only description" (General Discussion). Counter to this claim, the results of Experiment 1 show that in displays of simple geometrical shapes, participants benefitted from processing color-

modified descriptions relative to noun-only descriptions. As in the case of attending to all attributes in a description (rather than identifying the target as soon as possible), the results and predictions of Fukumura and Carminati (2021) may be specific to the displays of four objects of the same kind that they used in their study, and do not generalize to the more variable displays that have been used in the visual-world literature. If identifying a specific bow among four different bows requires a more demanding visual search than identifying a single bow in a display of four different objects, then the harder visual search could potentially mask the benefit of redundant modification—as the results of this study suggest.

Rather than relying on artificial displays of 2-D images, future studies should also investigate how redundant modification affects the listener's visual search for a referent in more realistic scenes. Rehrig et al. (under review) recently reported an eye-tracking study where participants benefitted from reading a redundant description of the target (e.g., "Find the black lamp" or "Find the lamp on the upper left") ahead of searching for it in a photograph of a real scene. While the results of Rehrig et al. support the visual efficiency hypothesis, future studies using naturalistic scenes should also measure visual search during language processing for a more accurate estimation of the potential benefits of redundant modification in real-time referent resolution (Engelhardt and Ferreira, 2016).

In conclusion, the results of this study confirm how perceptual discriminability affects both the production of redundant color adjectives and their efficiency for the listener's visual search, offering an answer to a 40-year debate: psycholinguistic studies reporting that people often use color words redundantly need not challenge the views from cognitive science and philosophy of language on efficient communication (Grice, 1975; Gibson et al., 2019). If we account for the psychophysics of referential communication and assume that a referential expression's efficiency is based not only on its informational value but also on its discriminatory value, then redundant color words can be more efficient than shorter descriptions, facilitating coordination between speakers and listeners.

Data availability

All data and analysis code are available from the Open Science Framework repository: https://osf.io/38m2h/?view_only=1292f98c00fa4d71abec0f1f91d38ff2.

Received: 18 October 2020; Accepted: 17 May 2021;

Published online: 17 June 2021

Notes

- 1 Thanks to Herb Clark for pointing out the importance of the restrictive / non-restrictive distinction in debates around referential over-specification (*personal communication*).
- 2 In line with the referential communication literature, I will continue to use the terms "redundant", "over-informative", and "over-specific" when qualifying adjectives that are not necessary to identify a referent in the visual context. However, under the visual efficiency hypothesis, those color adjectives are used descriptively (rather than restrictively) and are, therefore, not technically speaking redundant under the Gricean Maxim of Quantity.
- 3 In ongoing collaborative work, we are trying to identify the contribution of direct experience with visual search in reference comprehension on subsequent reference production, as potentially different from simple priming.
- 4 It is worth noting that, according to the visual efficiency hypothesis, pronominal expressions such as "the blue one" should be more efficient than minimal descriptions that do not include a color adjective (e.g., "the star"). The key difference between processing these two types of minimal referential expression is that searching by color should be faster than searching by shape (for eye-tracking evidence, see Rubio-Fernandez et al., 2020).

References

- Adams RC, Chambers CD (2012) Mapping the timecourse of goal-directed attention to location and colour in human vision. *Acta Psychologica* 139:515–523
- Arnold JE (2008) Reference production: production-internal and addressee-oriented processes. *Lang Cogn Process* 23:495–527
- Arts A, Maes A, Noordman L, Jansen C (2011) Over-specification facilitates object identification. *J Pragm* 43:361–374
- Bates D, Maechler M, Bolker B, Walker S (2014) lme4: linear mixed-effects models using Eigen and S4. <http://CRAN.R-project.org/package=lme4>.
- Belke E (2006) Visual determinants of preferred adjective order. *Vis Cogn* 14:261–294
- Belke E, Meyer AS (2002) Tracking the time course of multidimensional stimulus discrimination: analyses of viewing patterns and processing times during 'same'-different decisions. *Eur J Cogn Psychol* 14:237–266
- Bio BJ, Webb TW, Graziano MSA (2018) Projecting one's own spatial bias onto others during a theory-of-mind task. *Proc Natl Acad Sci USA*, 115: E1684–E1689
- Brainard DH, Freeman WT (1997) Bayesian color constancy. *J Opt Soc Am A* 14:1393–1411
- Bramão I, Reis A, Petersson KM, Faisca L (2011) The role of colour information on object recognition: a review and meta-analysis. *Acta Psychol* 138:244–253
- Buz E, Tanenhaus MK, Jaeger TF (2016) Dynamically adapted context-specific hyper-articulation: feedback from interlocutors affects speakers' subsequent pronunciations. *J Mem Lang* 89:68–86
- Christiansen MH (2019) Implicit-statistical learning: a tale of two literatures. *Top Cogn Sci* 11:468–481
- Clark HH, Marshall CR (1981) Definite reference and mutual knowledge. In: Joshi AK, Webber B, Sag IA (eds.) *Elements of discourse understanding*. Cambridge University Press, Cambridge, pp. 10–63
- Clark HH, Wilkes-Gibbs D (1986) Referring as a collaborative process. *Cognition* 22:1–39
- Clarke AD, Elsner M, Rohde H (2013) Where's Wally: the influence of visual salience on referring expression generation. *Front Psychol* 4:329
- Davidoff J (1991) *Cognition through color*. MIT Press, Cambridge, MA
- Davidoff J (2001) Language and perceptual categorisation. *Trend Cogn Sci* 5:382–387
- Davies C, Arnold JE (2019) Reference and informativeness: how context shapes referential choice. In: Cummins C, Katsos N (eds.) *Handbook of Experimental Semantics and Pragmatics*. Oxford University Press, Oxford, UK, pp. 474–493
- Davies C, Katsos N (2013) Are speakers and listeners 'only moderately Gricean'? An empirical response to Engelhardt et al. (2006) *J Pragmat* 49:78–106
- Davies C, Kreysa H (2017) Looking at a contrast object before speaking boosts referential informativeness, but is not essential. *Acta Psychol* 178:87–99
- Degen J, Hawkins RX, Graf C, Kreiss E, Goodman ND (2020) When redundancy is rational: a Bayesian approach to 'over-informative' referring expressions. *Psychol Rev* 127:591–621
- Dell GS, Brown PM (1991) Mechanisms for listener-adaptation in language production: limiting the role of the "model of the listener.". In: Napoli D, Kegl J (eds.) *Bridges between psychology and linguistics*. Academic Press, New York, NY, pp. 105–129
- Deschrijver E, Palmer C (2020) Reframing social cognition: relational versus representational mentalizing. *Psychol Bull* 146(11):941–969
- D'Zmura M, Lennie P (1986) Mechanisms of color constancy. *J Opt Soc Am A* 3:1662–1672
- Eberhard KM, Spivey-Knowlton MJ, Sedivy JC, Tanenhaus MK (1995) Eye movements as a window into real-time spoken language comprehension in natural contexts. *J Psychol Res* 24:409–436
- Engelhardt PE, Bailey KGD, Ferreira F (2006) Do speakers and listeners observe the Gricean Maxim of Quantity? *J Mem Lang* 54:554–573
- Engelhardt PE, Demiral ŞB, Ferreira F (2011) Over-specified referring expressions impair comprehension: an ERP study. *Brain Cogn* 77:304–314
- Engelhardt PE, Ferreira F (2016) Reaching sentence and reference meaning. In: Knoeferle P, Pykkönen-Klauck P, Crocker MW (eds.) *Visually situated language comprehension*. John Benjamins, pp. 93–127
- Ferreria VS (2019) A mechanistic framework for explaining audience design in language production. *Ann Rev Psychol* 70:29–51
- Frank MC, Goodman ND (2012) Predicting pragmatic reasoning in language games. *Science* 336:998
- Fukumura K, Carminati MN (2021) Over-specification and incremental referential processing: an eye-tracking study. *J Exp Psychol Learn Mem Cogn* (in press)
- Gatt A, Krahmer E, Van Deemter K, Van Gompel RP (2014) Models and empirical data for the production of referring expressions. *Lang Cogn Neurosci* 29(8):899–911
- Gatt A, Krahmer E, Van Deemter K, van Gompel RP (2017) Reference production as search: the impact of domain size on the production of distinguishing descriptions. *Cogn Sc* 41:1457–1492
- Gegenfurtner K, Rieger J (2000) Sensory and cognitive contributions of color to the recognition of natural scenes. *Curr Biol* 10:805–808

- Geurts B, Rubio-Fernandez P (2015) Pragmatics and processing. *Ratio* 28:446–469
- Gibson E, Futrell R, Piantadosi ST, Dautriche I, Mahowald K, Bergen L, Levy RP (2019) How efficiency shapes human language. *Trend Cogn Sci* 23:389–407
- Goodman ND, Frank MC (2016) Pragmatic language interpretation as probabilistic inference. *Trend Cogn Sci* 20:818. 829
- Graham NVS (2001) *Visual pattern analyzers*. Oxford University Press, Oxford
- Graziano MSA (2013) *Consciousness and the Social Brain*. Oxford University Press
- Graziano MSA, Kastner S (2011) Human consciousness and its relationship to social neuroscience: a novel hypothesis. *Cogn Neurosci* 2:98–113
- Grice HP (1975) Logic and conversation. In: Cole P, Morgan J (eds.) *Speech acts*. Academic Press, New York, NY, pp. 41–58
- Grodner D, Sedivy J (2011) The effect of speaker-specific information on pragmatic inferences. In: Pearlmutter N, Gibson E (eds.) *The processing and acquisition of reference*. MIT Press, Cambridge, MA, pp. 239–272
- Irwin DE, Colcombe AM, Kramer AF, Hahn S (2000) Attentional and oculomotor capture by onset, luminance and color singletons. *Vis Res* 40:1443–1458
- Jaeger TF, Ferreira V (2013) Seeking predictions from a predictive framework. *Behav Brain Sci* 36:359
- Jara-Ettinger J, Rubio-Fernandez P (2021) A computational model of reference production based on listener visual-search costs. Paper presented at the 34th CUNY Conference on Human Sentence Processing (CUNY2021). University of Pennsylvania, virtual meeting
- Jara-Ettinger J, Rubio-Fernandez P (under review). The social basis of referential communication: speakers construct reference based on listeners' expected visual search
- Keysar B, Barr DJ, Balin JA, Brauner JS (2000) Taking perspective in conversation: the role of mutual knowledge in comprehension. *Psychol Sci* 11:32–38
- Keysar B, Lin S, Barr DJ (2003) Limits on theory of mind use in adults. *Cognition* 89:25–41
- Koolen R, Goudbeek M, Krahmer E (2013) The effect of scene variation on the redundant use of colour in definite reference. *Cogn Sci* 37:395–411
- Koolen R, Krahmer E, Swerts M (2015) How distractor objects trigger referential over-specification: testing the effects of visual clutter and distractor distance. *Cogn Sci* 40:1617–1647
- Kurumada C, Jaeger TF (2015) Communicative efficiency in language production: optional case-marking in Japanese. *J Mem Lang* 83:152–178
- Krahmer E, Van Deemter K (2012) Computational generation of referring expressions: a survey. *Comput Linguist* 38(1):173–218
- Landau B, Gleitman LR (2009) *Language and experience: evidence from the blind child* (Vol. 8). Harvard University Press, Harvard
- Long M, Horton WS, Rohde H, Sorace A (2018) Individual differences in switching and inhibition predict perspective-taking across the lifespan. *Cognition* 170:25–30
- Long M, Moore I, Mollica F, Rubio-Fernandez P (under review). Perceptual contrast as a visual heuristic in the production of referential expressions
- Long M, Rohde H, Rubio-Fernandez P (2020) The pressure to communicate efficiently continues to shape language use later in life. *Sci Rep* 10(1):1–13
- Lu ZL, Doshier B (2013) *Visual psychophysics: from laboratory to theory*. MIT Press
- Lupyan G, Spivey MJ (2010) Redundant spoken labels facilitate perception of multiple items. *Atten Percept Psychophys* 72:2236–2253
- Maes A, Arts A, Noordman L (2004) Reference management in instructive discourse. *Discourse Process* 37:117–144
- Maloney LT, Wandell BA (1986) Color constancy: a method for recovering surface spectral reflectance. *J Opt Soc Am A* 3:29–33
- Mangold R, Pobel R (1988) Informativeness and instrumentality in referential communication. *J Lang Soc Psychol* 7:181–191
- Marr D (1982) *Vision: a computational investigation into the human representation and processing of visual information*. W.H. Freeman, San Francisco
- Moll H, Carpenter M, Tomasello M (2007) Fourteen-month-olds know what others experience only in joint engagement. *Dev Sci* 10(6):826–835
- Moll H, Tomasello M (2006) Level 1 perspective-taking at 24 months of age. *Br J Dev Psychol* 24(3):603–613
- Paraboni I, van Deemter K (2014) Reference and the facilitation of search in spatial domains. *Lang Cogn Neurosci* 29:1002–1017
- Paraboni I, Van Deemter K, Masthoff J (2007) Generating referring expressions: making referents easy to identify. *Comput Linguist* 33:229–254
- Pechmann T (1989) Incremental speech production and referential over-specification. *Linguistics* 27:89–110
- Pickering MJ, Garrod S (2004) Toward a mechanistic psychology of dialogue. *Behav Brain Sci* 27(2):169–190
- R Core Team (2018) *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria, <https://www.R-project.org/>
- Realí F, Spivey MJ, Tyler MJ, Terranova J (2006) Inefficient conjunction search made efficient by concurrent spoken delivery of target identity. *Percept Psychophys* 68:959–974
- Rehrig G, Cullimore RA, Henderson JM, Ferreira F (2021). When more is more: redundant modifiers can facilitate visual search. *Cognitive Research: Principles and Implications*, 6
- Rubio-Fernandez P (2016) How redundant are redundant colour adjectives? An efficiency-based analysis of colour over-specification. *Front Psychol*, 7
- Rubio-Fernandez P (2019) Over-informative speakers are cooperative: revisiting the Gricean Maxim of Quantity. *Cognit Sci*, 43
- Rubio-Fernandez P, Mollica F, Jara-Ettinger J (2020) Speakers and listeners exploit word order for communicative efficiency: a cross-linguistic investigation. *J Exp Psychol* 150:583–594
- Rubio-Fernandez P, Jara-Ettinger J (2020) Incrementality and efficiency shape pragmatics across languages. *Proceedings of the National Academy of Sciences* 117:13399–13404
- Ruby P, Decety J (2001) Effect of subjective perspective taking during simulation of action: a PET investigation of agency. *Nat Neurosci* 4:546–550
- Sedivy J (2003) Pragmatic versus form-based accounts of referential contrast: evidence for effects of informativity expectations. *J Psycholinguist Res* 32:3–23
- Sedivy J (2004) Evaluating explanations for referential context effects: evidence for Gricean mechanisms in online language interpretation. In: Trueswell J, Tanenhaus M (eds.) *Approaches to studying world-situated language use*. MIT Press, Cambridge, MA, pp. 345–364
- Sonnenschein S, Whitehurst GJ (1982) The effects of redundant communications on the behavior of listeners: does a picture need a thousand words? *J Psycholinguist Res* 11:115–125
- Spivey MJ, Tyler MJ, Eberhard KM, Tanenhaus MK (2001) Linguistically mediated visual search. *Psychol Sci* 12:282–286
- Tanaka J, Weiskopf D, Williams P (2001) The role of color in high-level vision. *Trend Cogn Sci* 5:211–215
- Tanenhaus MK, Chambers CC, Hanna JE (2004) Referential domains in spoken language comprehension: using eye movements to bridge the product and action traditions. In: *The interface of language, vision, and action: eye movements and the visual world*. In: Henderson JM, Ferreira F (eds.) *Psychology Press*. pp. 251–28
- Tourtour EN, Delogu F, Sikos L, Crocker MW (2019) Rational over-specification in visually-situated comprehension and production. *J Cult Cogn Sci* 3:175–202
- Treisman AM, Gelade G (1980) A feature-integration theory of attention. *Cogn Psychol* 12:97–136
- Umbach C (2006) Non-restrictive modification and backgrounding. In: *Proceedings of the Ninth Symposium on Logic and Language* (vol. 152159). Hungarian Academy of Sciences, Budapest
- van Deemter K, Gatt A, Van Gompel RP, Krahmer E (2012) Toward a computational psycholinguistics of reference production. *Topic Cogn Sci* 4(2):166–183
- van Gompel RP, van Deemter K, Gatt A, Snoeren R, Krahmer EJ (2019) Conceptualization in reference production: probabilistic modeling and experimental testing. *Psychol Rev* 126:345. 373
- Vogels J, Howcroft DM, Tourtour E, Demberg V (2020) How speakers adapt object descriptions to listeners under load. *Language. Cogn Neurosci* 35:78–92
- Wardlow-Lane W, Ferreira VS (2008) Speaker-external versus speaker-internal forces on utterance form: do cognitive demands override threats to referential success?. *J Exp Psychol Learn Mem Cogn* 34:1466–1481
- Wu S, Gibson E (2021) Word order predicts cross-linguistic differences in the production of redundant color and number modifiers. *Cogn Sci*
- Zadeh LA (1965) Fuzzy sets. *Inf Control* 8:338–353

Acknowledgements

Thanks to Chris Cummins, Julian Jara-Ettinger, Madeleine Long, Frank Mollica, and Laura Schulz for helpful discussions. This research was supported by a Young Research Talent Grant (Research Council of Norway, 230718). The author gratefully acknowledges this funding.

Competing interests

The author declares no competing interests.

Additional information

Correspondence and requests for materials should be addressed to P.R-F.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021