Color interpretation is guided by informativity expectations, not by world knowledge about colors

Hannah Rohde^{a,d}, Paula Rubio-Fernandez^{b,c,d}

^aLinguistics & English Language, University of Edinburgh

^bPhilosophy, Classics, History of Art & Ideas, University of Oslo

^cBrain & Cognitive Sciences, Massachusetts Institute of Technology

^dBoth authors contributed equally to this work.

ARTICLE INFO

Keywords:

language comprehension

overspecification

color typicality

Bayes Rule

eye tracking

ABSTRACT

When people hear words for objects with prototypical colors (e.g., ‘banana’), they look at objects of the same color (e.g., lemon). However, the experimental record also shows that when people speak, they tend to *omit* prototypical colors, using color adjectives when it is informative (e.g., when referring to clothes). These findings yield an interesting prediction, which we tested here: People may look at yellow objects when hearing ‘banana’, but they should look *away* from bananas when hearing ‘yellow’. The results of an offline sentence-completion task (N=100) and an online eye-tracking task (N=41) confirmed that when presented with truncated color descriptions (e.g., ‘Click on the yellow...’), people anticipate clothing items rather than stereotypical fruits. A corpus analysis ruled out the possibility that this association between color and clothing arises from simple co-occurrence statistics. We conclude that comprehenders make linguistic predictions based not only on what they *know* about the world (e.g., which objects are yellow) but also on what speakers tend to *say* about the world (i.e., their informativity expectations).

1. Introduction

Finishing each other’s sentences is often taken as a sign of knowing each other well. However, interlocutors’ synchronization applies not only to people who are particularly close: decades of psycholinguistics research reveals that language processing is, by its very nature, *predictive* (Kuperberg and Jaeger, 2016). Language comprehenders demonstrate remarkable expectation-driven processing, as revealed by a range of methodologies, from priming (McNamara, 2005) and reading (Van Berkum, Brown, Zwitserlood, Kooijman and Hagoort, 2005) to brain responses (DeLong, Urbach and Kutas, 2005) and anticipatory looking (Kamide, Altmann and Haywood, 2003).

*Corresponding author: Hannah Rohde, Hannah.Rohde@ed.ac.uk

For this study’s pre-registrations, as well as materials, data and analysis code, see:

Pre-registration Experiment 2: <https://osf.io/djzwc>

Pre-registration Experiment 3: <https://osf.io/zcwvs>

All materials, data, analysis: https://osf.io/sja5b/?view_only=859567206e3c4c41983f1e6c9c464753

Comprehenders use prediction at all levels of linguistic structure, relying on dependencies between sounds (DeLong et al., 2005), words (Kamide et al., 2003), syntactic structures (Levy, 2008), and propositions (Rohde and Horton, 2014). At the word level, comprehenders are particularly adept at using their world knowledge to activate relevant concepts. For example, people predict that ‘the mechanic checked...’ is more like to be followed by ‘the breaks’ than ‘the spelling’, whereas they make the reverse prediction if the agent is ‘the journalist’ (Bicknell, Elman, Hare, McRae and Kutas, 2010).

However, in production, speakers do not necessarily use language to talk about predictable content. On the contrary, speakers are known to *omit* content that is prototypical of an event. What speakers favor instead is content that is worth mentioning (e.g., an atypical instrument like ‘sweeping with a tree branch’; Brown and Dell, 1987; Lockridge and Brennan, 2002). This pattern of omission aligns with information-theoretic accounts that link predictability to reduction (Aylett and Turk, 2004; Levy and Jaeger, 2007) and more broadly to expectations of speaker cooperativity (Grice, 1975; Sperber and Wilson, 1986). Here we investigate comprehenders’ sensitivity to this informativity bias: Do comprehenders use not only their knowledge of the world (i.e. what is typical), but also their knowledge of informativity (i.e. what is newsworthy) to anticipate upcoming words? We focused on color words as a test case, and used truncated instructions to probe comprehenders’ expectations. Consider (1).

(1) Can you pass me the yellow ...

To anticipate the next word, comprehenders may rely on their knowledge about the color of objects in the world: Eye-tracking studies show that comprehenders look at yellow things after hearing a word for a typically yellow object (e.g., a lemon when hearing ‘banana’; Huettig and Altmann, 2011; Naor-Raz, Tarr and Kersten, 2003). Then again, if comprehenders assume the speaker is being informative, they may anticipate that the speaker is talking about an object that has low probability of possessing that color in the real world. This yields an interesting empirical prediction: Comprehenders may look at yellow objects when hearing ‘banana’, but they should look *away* from bananas when hearing ‘yellow’.

2. Color comprehension and production

The above eye-tracking studies suggest that the mention of an object triggers the activation of the corresponding color. However, not all objects are strongly associated to a specific color: some objects have prototypical colors (e.g., bananas), whereas others have variable colors (e.g., mugs). The con-

ceptual link between objects and colors is known to influence how speakers use color words. Color adjectives are often used *contrastively* to preempt an ambiguity between objects of the same kind (known as *competitors*; e.g., two bananas or two mugs), but they are also used descriptively (so-called *redundant* or *overspecific* uses (Pechmann, 1989)). Redundant color modifiers, however, are not produced at the same rate for all objects: Bananas are rarely described as yellow, for example, whereas non-prototypical colors are overspecified more (e.g., a yellow mug; Sedivy, 2003; Rubio-Fernandez, 2016; Westerbeek, Koolen and Maes, 2015).

Are comprehenders aware of the link between color typicality and production likelihood? Eye-tracking studies on *contrastive inference* suggest that they are: Comprehenders can reason that the use of color is contrastive, anticipating that the instruction ‘Pick up the yellow...’ will refer to the yellow object that has a competitor in the display, rather than to a yellow singleton (Sedivy, Tanenhaus, Chambers and Carlson, 1999; Sedivy, 2003). However, the strength of this contrastive inference (measured by anticipatory looking) depends on color typicality: The instruction ‘Pick up the yellow...’ triggers anticipatory looks to a banana (a prototypically yellow object) over a notebook (a non-prototypically yellow object), when there is a green banana in the display. However, that instruction does not trigger anticipatory looks when neither yellow object has a prototypical color (e.g., a mug and a ruler; Sedivy, 2004; Rubio-Fernandez, Mollica and Jara-Ettinger, 2021). These results suggest that comprehenders expect prototypical colors to be used contrastively (why else would one mention that a banana is yellow?), whereas non-prototypical colors are expected to be used descriptively (Sedivy, 2003, 2004). In line with this eye-tracking literature, a recent study using an offline sentence-completion task showed that comprehenders also expect color adjectives to be used descriptively when color is atypical (e.g., yellow strawberry; Kreiss and Degen, 2020).¹

Recent language production studies have shown that even when objects do not have a prototypical color, speakers tend to overspecify color more when it is central or relevant to a category (e.g., clothing) than when it is not (e.g., artefacts, geometrical shapes; Rubio-Fernandez, 2016, 2019). Language comprehension studies have confirmed that descriptive color adjectives can facilitate comprehenders’ visual search for the referent (Paraboni, Van Deemter and Masthoff, 2007; Arts, Maes, Noordman and Jansen, 2011; Tourtouri, Delogu, Sikos and Crocker, 2019; Rubio-Fernandez, in press), suggesting they support efficient communication (Rubio-Fernandez, 2016, 2019; Long, Rohde and Rubio-

¹ Sedivy reports a control condition where neither the yellow banana nor the yellow notebook had a competitor. In that condition, color overspecification had a marginal effect in favor of the non-prototypically yellow object (Sedivy, 2003, Footnote 5, p.17), suggesting that comprehenders are potentially sensitive to the different production likelihoods of prototypical and non-prototypical colors.

Fernandez, 2020). However, the tendency to mention the color of clothes seems so strong as to override efficiency considerations: When presented with a monochrome display of clothes, speakers produced expressions like ‘the yellow shirt’ around 40% of the time (despite color being a useless visual cue; Rubio-Fernandez, 2016), whereas they never produced ‘the yellow triangle’ in a monochrome display of shapes, and did so 40% of the time in polychrome displays (where color was an efficient visual cue, despite being formally redundant; Rubio-Fernandez, 2019).

Sedivy (2003; 2004)’s comprehension findings and Rubio-Fernandez (2016; 2019)’s production data raise an interesting question, which we address here: Do comprehenders predict that a redundant color word will be used to refer to objects for which color is relevant, not to objects for which that color is prototypical? For example, in hearing ‘the yellow...’, do comprehenders anticipate ‘shirt’, given speakers’ frequent production of redundant color adjectives for clothing, rather than ‘banana’, despite the strong conceptual link between yellow and bananas? Such a finding would go beyond previous work on comprehenders’ ability to make linguistic predictions based on their conceptual knowledge (i.e. what they *know* about the world) to show an additional ability to track speakers’ informative use of language (i.e. what speakers tend to *say* about the world).

3. Modelling conceptual probability and speaker informativity

In hearing ‘Can you pass me the yellow...’, a comprehender’s guess about the next word could rely entirely on general world knowledge about things that are typically yellow (*World Knowledge Hypothesis*). This would amount to using a plausibility-driven prior, assigning higher values to objects that are prototypically yellow (e.g., bananas) than to those that are not (e.g., strawberries or shirts). Alternatively, a comprehender could consider speaker production likelihoods alongside the prior (*Informativity Hypothesis*), as in (2).

$$(2) \quad p(\text{referent} = \text{object}_i | \text{'yellow'}) \propto p(\text{object}_{i, \text{COLOR}=\text{yellow}}) * p(\text{'yellow'} | \text{object}_{i, \text{COLOR}=\text{yellow}})$$

The conditional probability on the left captures comprehenders’ interpretation problem: What object is the speaker talking about given that they mentioned ‘yellow’? Under the World Knowledge Hypothesis, that probability is estimated directly from the predictability of the situation: the only term on the right-hand side would be $p(\text{object}_{i, \text{COLOR}=\text{yellow}})$, the probability that the object_i is yellow. Under the Informativity Hypothesis, an additional factor is incorporated: Given a yellow object that the speaker

could be talking about, how likely is the speaker to mention its color?

Note that (2) is not specifically a model of language production; it is simply Bayes' Rule, a theorem that establishes that the conditional probability on the left is proportional to the combination of the prior and likelihood on the right (for other applications in linguistics, see Kehler and Rohde, 2013; Rohde and Kurumada, 2018; Franke and Jäger, 2016). If predictive language processing depends in part on comprehenders' ability to reverse engineer what utterance a speaker is likely to be producing, casting the interpretation problem as in (2) highlights the two factors which determine the probability that a speaker is talking about a particular object. These two terms are at odds in 'Can you pass me the yellow...'. When estimating what object the speaker is requesting, a yellow strawberry is improbable (low situation prior) but if there were a yellow strawberry, that is the type of newsworthy situation that a speaker might be likely to say something about (high production likelihood) (for related models, see Kreiss and Degen, 2020; Rohde, Futrell and Lucas, 2021).

This intuition underlies the studies we present here. Given the ample experimental evidence that comprehenders use plausibility-driven priors during interpretation (Warren and Dickey, 2021), we ask whether they also show evidence of tracking and using production likelihoods. We first established whether our predictions could result from comprehenders simply keeping track of co-occurrence statistics, regardless of whether color is used contrastively or redundantly. Experiment 1 investigated that possibility in two analyses of the Google ngram corpus (Brants, Franz et al., 2006). Then, we conducted two comprehension studies: a web-based offline experiment, and an online eye-tracking experiment. In both tasks, participants indicated which word they anticipated after a truncated instruction like 'Click on the yellow...' when presented with a picture of prototypically yellow foods (e.g., two bananas) and a picture of non-inherently yellow clothes (e.g., two shirts or two skirts). We manipulated the redundant word in the instruction to compare responses to a color adjective ('yellow') versus a baseline number ('two'). In addition to this within-subject Instruction manipulation, a between-subjects manipulation varied the assignment of non-prototypical colors to objects (see Figure 1). This Pairing manipulation was included in Experiment 2 (and extended in Experiment 3 with the use of new pictures) to ensure that our findings were not attributable to the specific combinations of foods and clothes, or colors and clothes. Finally, we employed eye-tracking in Experiment 3 to investigate the possibility that early in the processing of color words, comprehenders may show an initial bias towards the prototypically colored foods, potentially resulting from a bottom-up effect of world knowledge (e.g., a strong association between yellow and bananas). Experiment 3 thus tests whether

world knowledge and informativity expectations inform language prediction at different points during processing.

4. Experiment 1: Corpus analysis of co-occurrence statistics

A pattern whereby color words bias participants towards anticipating a clothing word is predicted under the Informativity Hypothesis and is based on the assumption that comprehenders are aware that speakers produce redundant color adjectives at higher rates for objects that lack a prototypical color than for those with a prototypical color, especially for clothes, which are often described by color. However, it is also possible that a clothing bias could emerge if color+clothing collocations are simply more frequent in comprehenders' linguistic input and if comprehenders track those frequencies, independent of whether color is used contrastively or redundantly for clothes. In other words, rather than tracking speakers' descriptions of particular object categories (e.g., clothes vs foods) in the kinds of contexts we use here (i.e., contexts that lack a competitor object for which color might be discriminating), comprehenders could simply track the frequency of the phrases 'the yellow shirt' versus 'the yellow banana' in a context-independent way.

To investigate this possibility, we extracted corpus counts from the Google ngram corpus (Brants et al., 2006), whose contexts of use are unknown but whose ngram counts nonetheless provide an estimate of word collocation frequencies. We use the Google ngram corpus because it provides counts of approximately 1 trillion words from publicly-accessible web pages containing English text. The ngrams we target correspond to the Experiment 2 linguistic materials (see Supplemental Information) as well as generalised versions with a larger range of color~object and number~object collocations.

The goals of this corpus analysis are twofold: First to compare how often clothing versus food words appear in the trigrams that comprise our linguistic materials and, second, to compute Point-wise Mutual Information (PMI), an information theoretic measure that takes into account the relative frequencies of the words involved. The first metric used raw frequencies; we compared the trigram counts for clothing and food words in each of our experimental items in the color condition ('the yellow shirts' versus 'the yellow bananas') and the number condition ('the two shirts' versus 'the two bananas'). The second metric incorporated the relative frequency of the words themselves (to account for the fact that clothing words are simply more frequent overall than food words in the Google ngram corpus); we computed Point-wise Mutual Information (PMI) to test how close the association is between color and food words and between color and clothing words, and likewise between number and

Trigram	Experimental item usage	Clothing	Food
'the [color] [object]s'	Pairing 1 colors	69.8	355.5
'the [color] [object]s'	Pairing 2 colors	176.8	355.5
'the two [object]s'	Pairing 1 & Pairing 2	425.7	94.4
'the [color] [object]s'	generalized across all 10 colors	7902.0	2156.6
'the [number] [object]s'	generalized across all numbers	1457.9	189.3

Table 1

Mean frequencies for Experiment 2 linguistic materials, extracted from Google trigram counts where [color] is one of 10 colors and [object] is one of the 20 clothing and food objects in Table A.1. The first three rows show mean frequencies for the specific expressions in our experimental materials (e.g., 'the yellow shirts', 'the two shirts'). The last two rows show frequencies generalized across color and number. For example, the value 7902.0 represents the by-item mean of the summed raw trigram counts over all 10 colors for each clothing item (e.g., 'the yellow shirts', 'the red shirts', 'the white shirts', etc.).

food and between number and clothes (for a related analysis, see Culbertson, Schouwstra and Kirby, 2020). To foreshadow our findings, neither the trigram analysis nor the PMI analysis showed evidence that the pattern predicted under the Informativity Hypothesis could emerge from context-independent frequencies in comprehenders' linguistic input for the expressions used in our experimental items.

All the data and analysis scripts for the study are available at an Open Science Framework repository (see https://osf.io/sja5b/?view_only=859567206e3c4c41983f1e6c9c464753).

4.1. Trigram counts

Table 1 shows the mean counts for the trigram sequences in the Experiment 2 materials ('the two shirts', 'the two bananas', 'the yellow shirts', 'the yellow bananas', etc.). All counts represent the sum over lowercase and uppercase variants. The mean counts in the first three rows of Table 1 were computed as by-object means for the trigrams 'the two [object]s' and 'the [color] [object]s' for all objects in the Experiment 2 materials with the specific color object pairings that we used (e.g., 'the yellow shirts' and 'the yellow sandals' but not 'the orange shirts'). For the baseline condition 'the two [object]s', trigrams mentioning number and clothing are more frequent (425.7) than trigrams mentioning number and food (94.4), whereas for the critical color condition 'the [color] [object]s', this pattern is reversed such that the color words are followed more often by food words (355.5) than clothing words (69.8 for Pairing 1; 176.8 for Pairing 2). If comprehenders track frequencies like those in Table 1 and use them to anticipate upcoming words, we would expect to see behavior in keeping with the World Knowledge Hypothesis because our specific color words are followed more often by food words ('the yellow bananas') than clothing words ('the yellow shirts').

However, many of the specific expressions in Experiment 2 yielded a trigram count of 0. We therefore also report the mean counts computed over all possible color~object and number~object pairings.

The counts in the last two rows of Table 1 show that in both the color frame and the number frame, trigrams mentioning clothing are more frequent. If comprehenders track frequencies like these and use them to anticipate upcoming words, we would expect to see a clothing bias in both the color and number conditions, a pattern that matches neither the World Knowledge Hypothesis nor the Informativity Hypothesis. One should be cautious in interpreting the values in Table 1, however, because these raw frequencies are not adjusted for the relative frequency of individual words in the Google ngram corpus. For example, clothing words are simply more frequent than food words in the corpus, a pattern which could explain the high counts for clothing in the last two rows of Table 1.

4.2. Point-wise mutual information

To incorporate the relative frequency of the words in our critical expressions, we also compute PMI scores which measure the association between a pair of words: e.g., between the two words in ‘yellow shirts’ or between those in ‘yellow bananas’. PMI scores allow us to estimate whether the word ‘yellow’ appears more frequently with ‘shirts’ or with ‘bananas’, taking into account the frequency of the individual words. PMI uses the ratio of the probability of a pair of words appearing together and the independent probabilities of each of those words appearing at all. The PMI calculation is defined as in (i), illustrated with the words ‘yellow shirts’. We estimate PMI as in (ii) using a very large N of 1 trillion.

$$(i) \text{ PMI of the bigram 'yellow shirts' } = \log \frac{p(\text{'yellow shirts'})}{p(\text{'yellow'})p(\text{'shirts'})}$$

$$(ii) \text{ Estimated PMI of the bigram 'yellow shirts' } = \log \frac{\frac{\text{count}(\text{'yellow shirts'})}{N}}{\frac{\text{count}(\text{'yellow'})}{N} * \frac{\text{count}(\text{'shirts'})}{N}}$$

Table 2 reports the PMI scores for the specific bigrams in Experiment 2 (see first three rows). Those PMI scores show that the strength of association is higher between food words and their prototypical colors than between clothing words and the colors in which clothing was depicted in our materials. This presumably reflects the tight link between foods and their prototypical colors and the lack of a tight link between clothing and a specific color. The word ‘two’ in our materials is also more closely

Bigram	Experimental item usage	Clothing	Food
'[color] [object]s'	Pairing 1 colors	1.14	3.15
'[color] [object]s'	Pairing 2 colors	1.12	3.15
'two <objects>'	Pairing 1 & Pairing 2	0.51	1.26
'[color] [object]s'	generalized across 10 colors	1.45	1.40
'[number] [object]s'	generalized across numbers 0-10	-0.84	-0.08

Table 2

Point-wise Mutual Information for Experiment 2 linguistic materials, using Google bigram and unigram frequencies. The first three rows show PMIs for the specific bigrams in our experimental materials. The last two rows show PMIs generalized across color and number.

associated with food words than with clothing words, perhaps reflecting a property of foods (that they are counted and/or described with a low number) that differs from clothing (which are not necessarily quantified in pairs).

If comprehenders track the PMI scores of the Pairing 1 and Pairing 2 colors and use them to anticipate upcoming words, we would expect to see behavior in keeping with the World Knowledge Hypothesis because the PMI scores indicate that color words are more closely associated with food than clothing. Note that the PMI values for bigrams that contain colors words are higher than for bigrams that contain 'two', a finding that is in keeping with prior work that likewise reports a tighter association between objects and their properties (like color) than between objects and numbers (Culbertson et al., 2020).

Table 2 also reports the PMI values for color~object and number~object bigrams, generalized across colors and numbers. For those, we use bigrams containing any of the 10 color words in our experimental materials and any of the numbers 0-10 (as digits or words, excluding 1/'one' since the object noun was always plural). Taking the generalized color results first, the generalized color~object PMI scores reveal that clothing words appear with *any* color more than food words appear with *any* color, which makes sense if one considers that foods are associated with a specific prototypical color whereas clothing presumably can be linked to a range of colors. Participants in Experiment 2 were presented with color~object combinations from either Pairing 1 or Pairing 2, so the generalized PMI scores do not necessarily correspond to participant behavior when they encountered our specific items. That said, the PMI scores that generalise across color words raise the possibility that encountering *any* color word may bias comprehenders in favor of an upcoming clothing word, in keeping with the Informativity Hypothesis. For the PMI scores to predict participant behavior in this way, comprehenders would need to keep track of how speakers use the *class* of color words, beyond any specific word pairing in their input. Again, this is in line with the premise of the Informativity Hypothesis whereby

comprehenders track the rate at which speakers generally use color descriptors for objects of different categories.

For the generalised number results, these scores are also suggestive. The PMI scores in the last row of Table 2 are both negative, indicating a dissociation between numbers and the objects in our experiment, but the value for number~clothing is more negative than the value for number~food. This would raise the possibility that hearing *any* number word should yield a stronger bias against an upcoming clothing word than an upcoming food word. In a context with only food and clothing images, the presence of a number word could favor the food, since it bears the less negative association. Again, for these scores to drive participant behavior, comprehenders would need to track the rate at which speakers generally use number words for objects of different categories.

In sum, this analysis reports corpus frequencies in order to establish whether the behavior predicted under the Informativity Hypothesis could arise via simpler word associations in comprehenders' input. For our specific experimental materials, neither the trigram nor PMI measures correspond to the predicted pattern whereby color words favor clothing over food more so than number words. When we consider the generalised PMI scores, we do see a pattern that is in keeping with the Informativity hypothesis—namely that speakers use the *class* of color words more for clothing than for food. We conclude that behavior that accords with the Informativity Hypothesis would require an awareness of generalised speaker behavior, rather than specific associations between words in our experimental materials.

5. Experiment 2: Anticipation of object mention [pre-registered]

5.1. Methods

Participants

112 participants were recruited from the crowdsourcing platform Prolific and paid £3. Using Prolific's demographic selections, we recruited participants who were over 18, of US/UK Nationality, and monolingual English speakers. We excluded 5 participants who were not above chance on filler trials that had a correct response ($n=4$) or who did not complete the task ($n=8$). This yielded a dataset of 100 participants, 50 each for Pairings 1 and 2. Target participant numbers were estimated from an earlier pilot study.

Materials

Twenty experimental items were created using familiar foods and clothing items (see Appendix

Pairing

Instructions

[color] “Click on the yellow....”
[number] “Click on the two...”

Figure 1: Sample item with yellow objects. Pairing manipulation determines the clothing color by varying which clothing is presented with which food (banana~shirts versus banana~skirts). Instruction manipulation varies the type of redundant word (color versus number).

in Supplemental Information). The instructions for the experimental items contained 10 color words, each repeated twice within each Pairing. The food~clothing pairings were selected semi-randomly, with replacements when a particular color~clothing pairing corresponded to a brand name or fashion term that might make the expression more predictable (e.g., ‘Red Hat Linux’, ‘black tie’). For the experimental items, the position of the two pictures in each trial was counterbalanced (e.g., whether the bananas appeared on the left and the shirts on the right, or vice-versa). The color~food combinations were dictated by the foods themselves (e.g., purple eggplants, orange carrots). Counterbalancing of the visual presentation and the Instruction and Pairing manipulations resulted in 8 lists of materials.

The 20 experimental items were interleaved with an additional 35 filler items. Ten of these were similar to the experimental items with instructions that contained an ambiguous color word with no correct response (‘Click on the yellow...’), but the images were either both food (5) or both clothing (5). The remaining 15 fillers had a correct response cued by a color word (5), a number word (5), or an object category word (5). The position of the correct image was balanced across the filler types.

For all items, the instructions appeared as text above the images.

Procedure

Forced-choice responses were collected via a web-based interface that participants accessed from their own computer. Each item was presented on a page by itself with the two images displayed side by side. Participants were randomly assigned to one of the eight lists of materials. Trials were presented in one of a large set of possible pseudo-random orders (no two adjacent experimental trials).

Predictions

Under the World Knowledge Hypothesis, the color instruction should yield a food bias (compared to the baseline number condition) given the real-world link between foods and their prototypical colors. Under the Informativity Hypothesis, the color condition should instead yield a clothing bias, given the preference speakers show for producing redundant color adjectives for clothing items. Neither hypothesis predicts a main effect or interaction with Pairing since that manipulation is a control to vary the assignment of colors to clothing (i.e., the clothing images are the only images that change between Pairing 1 and Pairing 2).

5.2. Results

The binary outcome of clothing versus food was analysed using a logistic mixed effects model with the lme4 package (Bates, Kliegl, Vasishth and Baayen, 2015) in R (R Core Team, 2016). The model contained fixed effects for Instruction and Pairing and their interaction, and random effects for participants and items (where ‘item’ corresponded to the food image that was common across all Instructions and Pairings). Maximal random effect structure was used according to the experimental design, with random intercepts for participants and items, random by-item slopes for Pairing, and random by-participant and by-item slopes for Instruction. We centered the factors Instruction (number -0.5 , color $+0.5$) and Pairing (pairing1 -0.5 , pairing2, $+0.5$).

As expected under the Informativity Hypothesis, there was a main effect of Instruction with more clothing responses for the color instruction than for the number instruction ($\beta=0.853$, $SE=0.232$, $Z=3.671$, $p<0.001$). See Figure 2. There was no main effect of Pairing ($\beta=-0.364$, $SE=0.283$, $Z=1.288$, $p=0.20$) nor an Instruction $\times$ Pairing interaction ($\beta=0.168$, $SE=0.426$, $Z=0.394$, $p=0.69$), thereby confirming that participants’ click preferences were not driven by one particular set of pictures or assignment of colors to clothing.

6. Experiment 3: Eye-tracking anticipation of object mention [pre-registered analyses]

6.1. Methods

Participants

41 participants were recruited from the Subject Pool at MIT Brain and Cognitive Sciences Department and paid 10 USD. All participants were native speakers of English and reported having normal

Figure 2: Experiment 2 clothing responses by Pairing and Instruction.

color vision. Target participant numbers were estimated based on similar eye-tracking studies on color inferencing.

Materials

The same food-clothing pairs used in Experiment 2/Pairing 1 were used in Experiment 3. However, while the depicted objects were the same, the photographs were different to avoid unwanted effects of the specific photographs employed in the study. As in Experiment 2, the materials included 20 critical items and 35 fillers. However, in order to test real-time language processing, we presented the instructions in the auditory modality. The instructions were recorded by a female native speaker of American English. As in Experiment 2, instructions included color descriptions in the critical condition (e.g., ‘Click on the yellow...’) and a number description in the baseline condition (‘Click on the two...’). The position of the two images in each pair was counterbalanced and each item was matched with two different instructions (one including a color adjective and the other including the number ‘two’), which were fully crossed in 4 lists of materials. There was a 300ms preview window before the instructions started. Each pair of pictures was visible for 4000 ms (independent of participants’ clicks).

Procedure

Participants were asked to complete a series of truncated instructions by clicking on one of two images on the screen. Three measures were collected: responses (i.e. clicks on foods vs clothes),

response times (RTs; measured from the onset of the instructions) and proportion of fixation time on the two images (foods vs clothes).

Predictions

Experiment 3 tested the same predictions as Experiment 2: Under the World Knowledge Hypothesis, the color instruction should yield a food bias relative to the baseline number condition, whereas under the Informativity Hypothesis, the color condition should instead yield a clothing bias. In addition to replicating the response patterns observed in Experiment 2, we wanted to test these hypotheses in an eye-tracking paradigm in order to have online measures of language processing. In particular, we wanted to see whether in the color condition, participants' eye movements might reveal an initial food bias (driven by a bottom-up effect of world knowledge about colors) before they complete the instruction by selecting the clothes (as they did in Experiment 2). Such a pattern of results would suggest that world knowledge and informativity expectations (i.e., the prior and the likelihood terms in (2)) affect language prediction at different points during processing. Alternatively, participants' responses may be driven by their informativity expectations (i.e. redundant color adjectives are used to refer to clothes more often than to foods) from the earliest stages of processing, without suffering interference from the conceptual association between foods and prototypical colors (Huettig and Altmann, 2011; Naor-Raz et al., 2003). In addition to eye movements, RTs also tap real-time processing. For example, participants may click on clothes more often in the color condition than in the number condition, but may also respond more slowly. Such a pattern of results would also suggest that estimating color informativity is cognitively costly relative to a neutral baseline.

6.2. Results

Responses

As in Experiment 2, participants' binary choices between foods and clothes were analyzed using a logistic mixed-effects model with a fixed effect of Instruction (Color vs Number) and the maximal random effect structure for Participants and Items, including random intercepts for Participants and Items and random by-participant and by-item slopes for Instruction. Replicating the results of Experiment 2, and supporting the Informativity Hypothesis, there was a significant main effect of Instruction, with more clicks recorded on clothes in the Color condition than in the Number condition ($\beta = -1.067$, $SE = 0.297$, $Z = -3.590$, $p < 0.001$; see Figure 3).

Response Times

Figure 3: Experiment 3 clothing responses by Instruction (left) and RTs by object selected (right).

An analogous linear mixed-effects model was conducted on participants' RTs, but here the main effect of Instruction was not significant ($\beta=39.79$, $SE=32.79$, $df=32.63$, $t=1.214$, $p=0.234$), suggesting that participants made their picture selection at a comparable speed when they heard a color or a number description.

Following our pre-registration, we performed an exploratory RT analysis modelling the effect of Instruction and Response. That is, we investigated whether participants were faster or slower depending on what type of Instruction they heard (color vs number) and which picture they selected (clothes vs foods). For this analysis we used a linear mixed-effects model with fixed effects of Response and Instruction plus their interaction and the maximal random effect structure (because of model convergence issues, the interaction term between Instruction and Response was removed from the random effect structure of Item). The only significant result was a main effect of Response ($\beta=110.30$, $SE=50.34$, $df=41.01$, $t=2.191$, $p=0.034$), with faster RTs observed for clothes than for foods. The RT analysis does not suggest that participants' preference for clothes in the color condition is cognitively costly. Instead, they were generally faster to click on clothes than on foods.

Fixation time

Three time windows were pre-registered for the analysis of looking data: an early window corresponding to the critical word in the instructions (i.e. from the onset to the offset of the color adjective or the numeral, depending on the condition), an intermediate window ranging from the offset of the critical word until the participant's response, and a late window ranging from the participant's response

Figure 4: Experiment 3 proportion looks to clothing and food areas of interest in the Color condition (left) and Number condition (right). Solid vertical lines show the onset and average offset of the color/number words. Dashed vertical lines show the average click time.

until the end of the trial. Following the standard convention in visual world studies, time windows were adjusted by +200ms to account for the time it takes to launch a saccade. Separate analyses were performed for each time window (early, intermediate and late) for each instruction type (color vs number), as pre-registered. For visualizations of the proportion of fixation time during real-time processing, see Figure 4.

Early time window

Percentage of fixation times were analyzed using a linear mixed-effects model with a fixed effect of Area of Interest (AOI; clothes vs foods) and the maximal random effect structure for Participants and Items, including random intercepts for Participants and Items and random by-participant and by-item slopes for AOI. The model did not reveal significant results in the early time window, neither in the Number condition ($\beta=-1.177$, $SE=5.424$, $df=50.491$, $t=-0.217$, $p=0.829$), nor in the Color condition ($\beta=2.438$, $SE=4.497$, $df=49.831$, $t=0.542$, $p=0.59$). These results suggest that when comprehenders processed color adjectives, they did not have an early food bias driven by world knowledge, which was then overridden by an informativity bias. Instead, in the earliest stages of processing, participants did not show a preference for either clothes or foods.

Intermediate time window

The same model was constructed for the fixation-time data for the second time window. The model

revealed no significant results in the Number condition ($\beta=0.546$, $SE=4.987$, $df=54.464$, $t=0.109$, $p=0.913$), but a significant effect of AOI in the Color condition ($\beta=-16.763$, $SE=4.404$, $df=51.609$, $t=-3.806$, $p<0.001$), with participants fixating longer on the clothes than the foods. The results of the intermediate time window therefore confirm that starting from the offset of the color adjective, participants showed a preference for clothes driven by informativity expectations. There was therefore no evidence of an early stage in color word processing that is driven by world knowledge about color typicality.

Late time window

Analyses of the fixation data for the last window revealed non-significant results in the Color condition ($\beta=-4.023$, $SE=3.659$, $df=43.519$, $t=-1.099$, $p<0.278$), but a significant effect of AOI in the Number condition ($\beta=9.510$, $SE=3.857$, $df=47.221$, $t=2.465$, $p<0.018$), resulting from participants fixating more on the foods than the clothes once they had selected a picture. The Number condition was included as a baseline against which to compare behavior in the Color condition, for which the Informativity and World Knowledge Hypotheses make competing predictions. However, we speculate that participants' behavior in the Number condition —i.e., their late preference for foods over clothes— could reflect the flip side of those same predictions. If clothing carries an *association* with color words, then one would expect a concomitant *dissociation* with other words; this dissociation would align with the observed late bias against clothing for the Number condition and is also in keeping with the frequencies that characterise these associations/dissociations, as reported in Experiment 1.

7. General Discussion

Across two experiments, we tested comprehenders' expectations about upcoming words in contexts that pitted world knowledge against informativity expectations. Comprehenders consistently used a prenominal color adjective to anticipate the mention of clothing ('the yellow...shirts'), which lacks prototypical color. This color~clothing bias emerged even though the color word described the prototypical color of a depicted food. In offline picture selections and online eye movements, comprehenders showed no evidence of a color~food bias, contra models in which next-mention expectations are primarily driven by world knowledge ('the yellow... bananas'). Rather, the findings support our Informativity Hypothesis: Comprehenders use knowledge not only about the probability of different situations in the world, but also about speakers' use of language to describe the world.

Recent psycholinguistic studies have shown that comprehenders keep track of different sources of

distributional information in language (e.g., words, morphemes, phonemes) in speakers' productions, which is independent of their world knowledge. Arnon and Snider (2010) show that comprehenders are sensitive to the frequencies of compositional four-word phrases (e.g. 'Don't have to worry' vs 'Don't have to wait'), with faster processing for phrases that speakers produce with higher frequency. The authors conclude that language users not only keep track of word-level usage but also store phrase-frequency information spanning various levels of abstraction. More recently, Morgan and Levy (2016) investigated multi-word phrases whose semantic content is held constant. They show that online processing of highly frequent binomials (e.g., 'bread and butter') is primarily driven by direct experience (as estimated from corpus frequency counts), whereas online processing of novel binomial expressions (e.g., 'bishops and seamstresses' vs 'seamstresses and bishops') is influenced by abstract knowledge of the ordering constraints in speakers' productions (as estimated by a probabilistic model). The authors interpret their results as supporting models of language processing including both compositional generation and direct reuse of multi-word expressions. Our study contributes to this literature by showing sensitivity to another speaker production constraint: Alongside prior work showing that comprehenders can deploy their world knowledge of stereotypical colors in referential communication (e.g., looking at a lemon when hearing 'the banana...'; (Huettig and Altmann, 2011; Naor-Raz et al., 2003), our results show that comprehenders are sensitive to informativity constraints in the use of color adjectives in non-contrastive contexts. Comprehenders' behavior suggests an awareness of speakers' production preferences, namely their tendency to use color words redundantly more often for referents without prototypical colors, such as clothes.

The results of our eye-tracking experiment confirmed that informativity expectations constrain the interpretation of color adjectives early on in processing, rather than resulting from later, top-down processes inhibiting world knowledge associations (e.g., between the color yellow and bananas). By contrast, when processing a numerical description ('Click on the two...'), participants' eye movements revealed hesitation between the clothes and foods in the pictures, in line with their chance performance when selecting a sentence continuation. If the color condition had induced hesitation, we expect we would have seen a pattern that resembles that in the baseline numerical condition. These results contribute to the eye-tracking literature on adjective processing, which shows that comprehenders are sensitive to the frequency with which different types of adjectives are used contrastively or redundantly, deriving contrastive inferences when interpreting scalar and material adjectives, but not always with color adjectives (Sedivy, 2003, 2004; Sedivy et al., 1999; Aparicio, Xiang and Kennedy, 2016;

Rubio-Fernandez, Terrasa, Shukla and Jara-Ettinger, 2019). Our results further show that comprehenders are sensitive to the kinds of lexical categories that tend to be overspecified by color (at least when the choice is between clothes and foods). Future studies should investigate how fine-grained comprehenders' sensitivity to color overspecification may be depending on the lexical categories contrasted (e.g., when selecting between a white car and a white fridge).

Here we focus on color because its non-contrastive use is well-documented (Pechmann, 1989) and its inclusion/omission is strongly influenced by object category (Sedivy, 2003; Rubio-Fernandez, 2016; Westerbeek et al., 2015; Kreiss and Degen, 2020). However, the results should extend to descriptors with weaker color biases or those outside the visual domain. For example, both vehicles and appliances lack inherent color, but speakers mention color more for cars than for appliances. Likewise, some objects have inherent properties (old ruins, modern smartphones), and speakers may thus prefer to use those descriptors informatively for other objects (old/new cars, modern/antique furniture). If speakers reliably show such production biases and if comprehenders can effectively track them, new questions emerge regarding the domains where informativity expectations might arise.

8. Summary and conclusions

Ample empirical evidence shows that language comprehension is predictive in nature. For example, people predict that 'the mechanic checked...' is more like to be followed by 'the breaks' than by 'the spelling', whereas they make the reverse prediction if the agent is 'the journalist'. Here we use Bayes' Rule to formalize a related hypothesis: When language comprehenders anticipate a noun following a color adjective (e.g., 'Can you pass me the red...'), they rely not only on their world knowledge about colors (e.g., which objects are typically red), but also on their experience of speakers' use of color adjectives to refer to different objects (e.g., the frequency of 'the red strawberries' vs 'the red shoes'). Our findings go beyond previous work on people's ability to make linguistic predictions based on their world knowledge, and show how comprehenders can reverse engineer the way speakers tend to talk about the world.

9. Author Contributions

Both authors contributed equally to the study design, data collection, analysis, and manuscript drafting and revision.

10. Acknowledgements

This research was supported by a FRIPRO Grant from the Research Council of Norway (Ref. 275505) awarded to PRF and a Leverhulme Trust Prize in Languages and Literatures to HR (Ref. PLP-2016-128). The authors gratefully acknowledge this funding. Thanks also to Francis Mollica for help with data visualizations.

References

- Aparicio, H., Xiang, M., Kennedy, C., 2016. Processing gradable adjectives in context: A visual world study. *Semantics and Linguistic Theory* 25, 413–432.
- Arnon, I., Snider, N., 2010. More than words: Frequency effects for multi-word phrases. *Journal of Memory and Language* 62, 67–82.
- Arts, A., Maes, A., Noordman, L., Jansen, C., 2011. Overspecification facilitates object identification. *Journal of Pragmatics* 43, 361–374.
- Aylett, M., Turk, A., 2004. The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech. *Language and Speech* 47, 31–56.
- Bates, D., Kliegl, R., Vasishth, S., Baayen, H., 2015. Parsimonious mixed models. *arXiv preprint arXiv:1506.04967*.
- Bicknell, K., Elman, J.L., Hare, M., McRae, K., Kutas, M., 2010. Effects of event knowledge in processing verbal arguments. *Journal of Memory and Language* 63, 489–505.
- Brants, T., Franz, A., et al., 2006. Web it 5-gram version 1. Philadelphia: Linguistic Data Consortium.
- Brown, P.M., Dell, G.S., 1987. Adapting production to comprehension: The explicit mention of instruments. *Cognitive Psychology* 19, 441–472.
- Culbertson, J., Schouwstra, M., Kirby, S., 2020. From the world to word order: deriving biases in noun phrase order from statistical properties of the world. *Language* 96, 696–717.
- DeLong, K.A., Urbach, T.P., Kutas, M., 2005. Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience* 8, 1117–1121.
- Franke, M., Jäger, G., 2016. Probabilistic pragmatics, or why bayes' rule is probably important for pragmatics. *Zeitschrift für Sprachwissenschaft* 35, 3–44.
- Grice, H.P., 1975. Logic and conversation, in: *Speech Acts*. Brill, pp. 41–58.
- Huetig, F., Altmann, G.T., 2011. Looking at anything that is green when hearing “frog”: How object surface colour and stored object colour knowledge influence language-mediated overt attention. *Quarterly Journal of Experimental Psychology* 64, 122–145.
- Kamide, Y., Altmann, G.T., Haywood, S.L., 2003. The time-course of prediction in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language* 49, 133–156.
- Kehler, A., Rohde, H., 2013. A probabilistic reconciliation of coherence-driven and centering-driven theories of pronoun interpretation. *Theoretical Linguistics* 39, 1–37.

- Kreiss, E., Degen, J., 2020. Production expectations modulate contrastive inference, in: *Proceedings of the 42nd Annual Meeting of the Cognitive Science Society*, pp. 259–265.
- Kuperberg, G.R., Jaeger, T.F., 2016. What do we mean by prediction in language comprehension? *Language, Cognition and Neuroscience* 31, 32–59.
- Levy, R., 2008. Expectation-based syntactic comprehension. *Cognition* 106, 1126–1177.
- Levy, R., Jaeger, T.F., 2007. Speakers optimize information density through syntactic reduction. *Advances in Neural Information Processing Systems* 19, 849.
- Lockridge, C.B., Brennan, S.E., 2002. Addressees' needs influence speakers' early syntactic choices. *Psychonomic bulletin & review* 9, 550–557.
- Long, M., Rohde, H., Rubio-Fernandez, P., 2020. The pressure to communicate efficiently continues to shape language use later in life. *Scientific Reports* 10, 1–13.
- McNamara, T.P., 2005. *Semantic priming: Perspectives from memory and word recognition*. Psychology Press.
- Morgan, E., Levy, R., 2016. Abstract knowledge versus direct experience in processing of binomial expressions. *Cognition* 157, 384–402.
- Naor-Raz, G., Tarr, M.J., Kersten, D., 2003. Is color an intrinsic property of object representation? *Perception* 32, 667–680.
- Paraboni, I., Van Deemter, K., Masthoff, J., 2007. Generating referring expressions: Making referents easy to identify. *Computational Linguistics* 33, 229–254.
- Pechmann, T., 1989. Incremental speech production and referential overspecification. *Linguistics* 27, 89–110.
- R Core Team, 2016. *R: A language and environment for statistical computing*.
- Rohde, H., Futrell, R., Lucas, C.G., 2021. What's new? a comprehension bias in favor of informativity. *Cognition* 209, 104491.
- Rohde, H., Horton, W.S., 2014. Anticipatory looks reveal expectations about discourse relations. *Cognition* 133, 667–691.
- Rohde, H., Kurumada, C., 2018. Alternatives and inferences in the communication of meaning, in: *Psychology of Learning and Motivation*. Elsevier. volume 68, pp. 215–261.
- Rubio-Fernandez, P., 2016. How redundant are redundant color adjectives? an efficiency-based analysis of color overspecification. *Frontiers in Psychology* 7, 153.
- Rubio-Fernandez, P., 2019. Overinformative speakers are cooperative: Revisiting the gricean maxim of quantity. *Cognitive Science* 43, e12797.
- Rubio-Fernandez, P., in press. Color discriminability makes over-specification efficient: Theoretical analysis and empirical evidence. *Humanities and Social Sciences Communications*.
- Rubio-Fernandez, P., Mollica, F., Jara-Ettinger, J., 2021. Speakers and listeners exploit word order for communicative efficiency: A cross-linguistic investigation. *Journal of Experimental Psychology: General* 150, 583–594.
- Rubio-Fernandez, P., Terrasa, H.A., Shukla, V., Jara-Ettinger, J., 2019. Contrastive inferences are sensitive to informativity expectations, adjective semantics and visual salience URL: psyarxiv.com/mr4ah, doi:10.31234/osf.io/mr4ah.
- Sedivy, J.C., 2003. Pragmatic versus form-based accounts of referential contrast: Evidence for effects of informativity expectations. *Journal of Psycholinguistic Research* 32, 3–23.
- Sedivy, J.C., 2004. Evaluating explanations for referential context effects: Evidence for gricean mechanisms in online language interpretation, in: *Approaches to studying world-situated language use: Bridging the language-as-product and language-as-*

- action traditions. MIT Press, pp. 345–364.
- Sedivy, J.C., Tanenhaus, M.K., Chambers, C.G., Carlson, G.N., 1999. Achieving incremental semantic interpretation through contextual representation. *Cognition* 71, 109–147.
- Sperber, D., Wilson, D., 1986. *Relevance: Communication and cognition*. volume 142. Harvard University Press Cambridge, MA.
- Tourtouri, E.N., Delogu, F., Sikos, L., Crocker, M.W., 2019. Rational over-specification in visually-situated comprehension and production. *Journal of Cultural Cognitive Science* 3, 175–202.
- Van Berkum, J.J., Brown, C.M., Zwitserlood, P., Kooijman, V., Hagoort, P., 2005. Anticipating upcoming words in discourse: evidence from erps and reading times. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31, 443.
- Warren, T., Dickey, M.W., 2021. The use of linguistic and world knowledge in language processing. *Language and Linguistics Compass* 15, e12411.
- Westerbeek, H., Koolen, R., Maes, A., 2015. Stored object knowledge and the production of referring expressions: The case of color typicality. *Frontiers in Psychology* 6, 935.

Instruction	Pairing 1	Pairing 2
'Click on the black/two...'	beans, aprons	beans, slippers
'Click on the black/two...'	grapes, sweatshirts	beans, swimsuits
'Click on the blue/two...'	plums, ties	plums, aprons
'Click on the blue/two...'	bluecorn, swimsuits	corn, sweatshirts
'Click on the brown/two...'	coconuts, gloves	coconuts, dresses
'Click on the brown/two...'	onions, slippers	onions, ties
'Click on the green/two...'	cucumbers, dresses	cucumbers, gloves
'Click on the green/two...'	zucchini, boots	zucchini, hats
'Click on the orange/two...'	carrots, hats	carrots, boots
'Click on the orange/two...'	pumpkins, coats	pumpkins, blouses
'Click on the pink/two...'	guavas, blouses	guavas, jackets
'Click on the pink/two...'	sausages, belts	sausages, coats
'Click on the purple/two...'	eggplants, jackets	eggplants, belts
'Click on the purple/two...'	figs, high heels	figs, scarves
'Click on the red/two...'	cherries, scarves	cherries, high heels
'Click on the red/two...'	tomatoes, sweaters	tomatoes, bras
'Click on the white/two...'	garlic, bras	garlic, shirts
'Click on the white/two...'	marshmallows, skirts	marshmallows, sandals
'Click on the yellow/two...'	lemons, sandals	lemons, sweaters
'Click on the yellow/two...'	bananas, shirts	bananas, skirts

Table A.1

Materials for Experiment 2 (Pairing 1 and Pairing 2) and Experiment 3 (Pairing 1).

Supplementary Information

Table A.1 shows the food and clothing items that were paired in Experiments 2 and 3.