

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Color Overmodification Emerges from Data-Driven Learning and Pragmatic Reasoning

Permalink

<https://escholarship.org/uc/item/9kn7n6qb>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Fang, Fei
Sinha, Kunal
Goodman, Noah
et al.

Publication Date

2022

Peer reviewed

Color Overmodification Emerges from Data-Driven Learning and Pragmatic Reasoning

Fei Fang¹

Kunal Sinha²

Noah D. Goodman^{3,1}

Christopher Potts⁴

Elisa Kreiss⁴

¹Department of Computer Science ²Department of Symbolic Systems

³Department of Psychology ⁴Department of Linguistics

Stanford University

Stanford, CA 94305 USA

{feifang, ksinha2, ngoodman, cgpotts, ekreiss}@stanford.edu

Abstract

Speakers’ referential expressions often depart from communicative ideals in ways that help illuminate the nature of pragmatic language use. Patterns of *overmodification*, in which a speaker uses a modifier that is redundant given their communicative goal, have proven especially informative in this regard. It seems likely that these patterns are shaped by the environment a speaker is exposed to in complex ways. Unfortunately, systematically manipulating these factors during human language acquisition is impossible. In this paper, we propose to address this limitation by adopting neural networks (NN) as learning agents. By systematically varying the environments in which these agents are trained, while keeping the NN architecture constant, we show that overmodification is more likely with environmental features that are infrequent or salient. We show that these findings emerge naturally in the context of a probabilistic model of pragmatic communication.

Keywords: overmodification; reference; pragmatics; learning semantics; rational speech acts

Introduction

Overmodification describes the well-attested tendency of speakers to supply more information than apparently needed for a listener to identify the intended referent. Consider the example in Figure 1 (top). The speaker’s goal is to communicate the target image (here highlighted with a yellow box) to a listener. While the utterance “circle” would suffice for the listener to pick out the target from the context, speakers often use “red circle” instead, whereas the unmodified noun “circle” is more likely when circles are typically red (middle) or when color in general is not as salient (bottom). In this paper, we show that such patterns arise naturally in neural network speaker agents that are (1) trained in these different environments and (2) reason pragmatically about a listener.

The rate of overmodification with color terms has been shown to vary in structured ways (Pechmann, 1989; Engelhardt et al., 2006; Koolen et al., 2011). Non-color-diagnostic objects, which occur in a variety of colors equally frequently (e.g., cups), are readily overmodified (Sedivy, 2003; Rubio-Fernández, 2016). For objects that are associated with one particular color (e.g., bananas), the overmodification rate drops when the color is typical (“banana” for a yellow banana), and increases when it is atypical (“blue banana” for a blue banana) (Westerbeek et al., 2015; Degen et al., 2020; Sedivy, 2003; Kreiss & Degen, 2020).

In addition to color diagnosticity, the visual salience of color has also been argued to play an important role in why

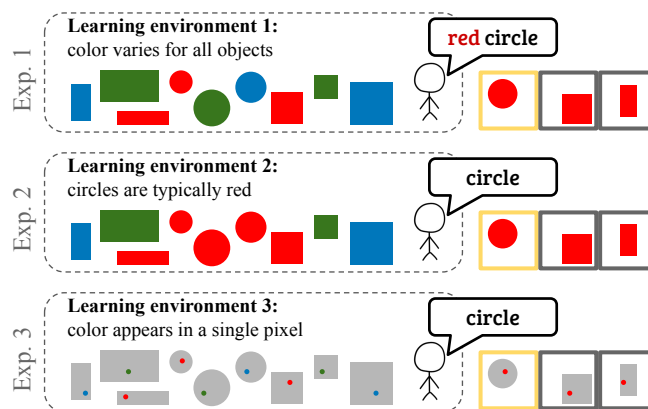


Figure 1: Neural speakers are exposed to different environments during learning and tested on their choice of color overmodification. Varying the underlying data distribution shapes neural speakers’ overinformative usage of color, explaining previous findings from human production and suggesting additional factors that may give rise to overmodification.

speakers tend to frequently overmodify with color but not with other modifiers such as size (Sedivy, 2003; van Gompel et al., 2019; Davies & Katsos, 2013; Rubio-Fernandez et al., 2019; Rubio-Fernández, 2016; Brown-Schmidt & Konopka, 2011). This intuition predicts that overmodification rates decrease with reduced visual salience.

A variety of computational models have been proposed to account for overmodification patterns. The Incremental algorithm (Dale & Reiter, 1995) predicts color overmodification whenever it serves to disambiguate from at least one distractor. The probabilistic referential overspecification model (PRO; van Gompel et al., 2019) predicts that speaker preferences and visual salience explain the overmodification patterns by arguing that atypically colored objects are more salient. Recently, several models of overmodification have been proposed within the Rational Speech Acts (RSA) framework (Frank & Goodman, 2012; Goodman & Frank, 2016), in which overmodification arises from abstract speaker agents reasoning about internal listeners (Cohn-Gordon et al., 2019; Waldon & Degen, 2021). Degen et al. (2020) propose an RSA-based model that predicts overmodification patterns by defining pragmatic reasoning over non-deterministic seman-

tics. In this way, the typicality effect, for instance, arises from the perception that the utterance “banana” applies less well to blue bananas than yellow ones.

What could lead a language learner to adopt the sort of semantics that produces such overmodification patterns? All previously proposed models assume that the speaker is already equipped with a complete semantics, abstracting away the process through which an agent extracts attributes of an object from raw visual input. These models therefore cannot provide a perspective on how learning to use language in the world might give rise to such a semantics in the first place. We overcome this limitation by introducing a neural network that learns the semantics from data consisting of raw visual and linguistic input. This allows us to show how overmodification arises directly from the underlying data distribution.

Our neural network-based approach also allows us to explore the role of pragmatic reasoning in overmodification. To do this, we evaluate a neural pragmatic speaker model (*RSA speaker*) and a neural literal speaker model (*literal speaker*) while varying the underlying data distributions they learn from. Both the literal speaker and the RSA speaker receive the context in the form of images as visual input, and do not observe any systematic overmodification during training. In a series of experiments, we manipulate the data that the models learn from by inducing color typicality (Exp. 2) as well as decreasing the salience of color (Exp. 3), and we compare these results to a baseline distribution (Exp. 1). We conclude that both pragmatic reasoning about a listener and properties of the underlying data distribution play a central role in achieving human-like rates of color overmodification. This provides a candidate explanation for the underlying processes giving rise to color overmodification in human language.

Background

Empirical Evidence for Overmodification

Overmodification by speakers is a robustly reported phenomenon on referring expression production (Pechmann, 1989; Nadig & Sedivy, 2002; Sedivy, 2003; Engelhardt et al., 2006; Koolen et al., 2011; Rubio-Fernández, 2016; Degen et al., 2020; Heller, 2020). Color adjectives stand out since their rate of overmodification is generally higher than for other modifiers such as size and material (Pechmann, 1989; Sedivy, 2005; Mitchell et al., 2013; Gatt et al., 2011). Two main sources have been proposed to explain this. First, in contrast to the relative nature of size adjectives, color adjectives are considered absolute since their attribution to an object is (relatively) context-independent (Kennedy, 2007; Syrett et al., 2009). Second, color is inherently visually salient and therefore prioritized in production (Sedivy, 2003; van Gompel et al., 2019; Davies & Katsos, 2013; Rubio-Fernandez et al., 2019; Rubio-Fernández, 2016; Brown-Schmidt & Konopka, 2011; Long et al., 2021; Rubio-Fernandez, 2021) due to speaker-internal or listener-oriented processes (Arnold, 2008).

Even within the domain of color adjectives, overmodifica-

tion rates vary based on how closely associated the object is with the presented color, i.e., the object’s color-diagnostics. Based on object recognition experiments, Tanaka and Presnell (1999) motivate a distinction of objects with high and low color-diagnostics. Low- or non-color-diagnostic objects can plausibly occur in many colors (e.g., cups), whereas highly color-diagnostic objects only have a few highly associated colors (e.g., bananas). While speakers often overmodify with color when referring to non-color-diagnostic objects, they rarely do so with color-diagnostic ones that are presented in their typical color (Sedivy, 2003). Moreover, the rate of overmodification has been shown to increase linearly with decreasing typicality of the color for the object (Westerbeek et al., 2015; Degen et al., 2020). For instance, a yellow banana would receive the least, a blue banana the most, and a brown banana an intermediate rate of overmodification.

While this literature rigorously explores color overmodification, the question of how changes in the visual environment of a learning agent causally drive overmodification remains elusive to empirical methods. In Exp. 1 and 2, we explore how the general phenomenon of overmodification arises naturally from learning from distinct data distributions. Since these data distributions can be arbitrarily adjusted, they also offer the opportunity to simulate speakers in worlds distinct from ours and study their overmodification patterns (Exp. 3).

Rational Speech Acts (RSA) framework

The Rational Speech Acts (RSA) framework is a family of computational models of communication (Frank & Goodman, 2012) where two rational agents, a speaker and a listener, are modeled as recursively reasoning about each other when producing and interpreting utterances. This work follows important precedents from game theory (Lewis, 1969; Jager, 2007; Franke, 2009; Golland et al., 2010), and has been shown to capture a wide range of linguistic phenomena (e.g., Goodman and Stuhlmüller, 2013; Kao et al., 2014; Hawkins and Goodman, 2017; van Tiel et al., 2021).

The probabilistic nature of the RSA framework lends itself naturally to integration with neural network systems. Recent work has shown that RSA reasoning helps in a variety of areas, including continual language adaptation (Hawkins & Goodman, 2017), color reference in context (Monroe et al., 2017, 2018), following navigational instructions (Fried et al., 2018), and context-aware image-based text generation (Andreas & Klein, 2016; Vedantam et al., 2017; Cohn-Gordon et al., 2018; Nie et al., 2020).

While previous work has compared the *overall performance* of neural RSA speakers versus non-RSA speakers (White et al., 2020; Andreas & Klein, 2016; Monroe et al., 2017), we specifically focus on the extent to which such agents rely on overmodification (specifically with color) to communicate. We further extend previous work by investigating how changes in the underlying data distribution that the models learn from affects pragmatic language use. This leads to an account in which overmodification emerges naturally from the visual and distributional properties of the world

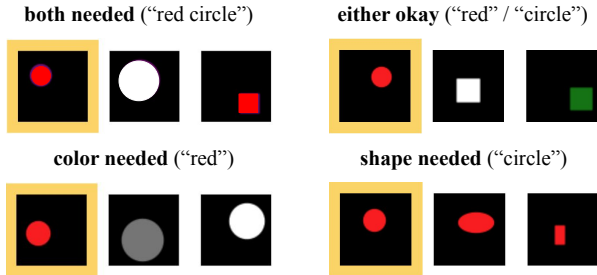


Figure 2: Context conditions, varying in the type of information that a referring expression must provide to distinguish the target. Color overmodification occurs when color is mentioned in the *shape needed* condition.

that a pragmatic speaker is exposed to during learning.

Task

We explore the production of referring expressions in simple reference games in which a speaker and listener are presented with the same set of three images, varying in order. The speaker’s goal is to communicate the identity of the target image to the listener, whose goal is to pick out the target.

The games are generated synthetically using the Shape-World dataset creation framework (Kuhnle & Copestake, 2017). In this setup, six colors are randomly paired with four potential shapes, resulting in 24 unique color–shape compositions. Each generated image displays one colored object that is sized and positioned randomly against a black background. Each generated game consists of three distinct images and varies in the minimal amount of information the listener would need to single out the target, resulting in the four *context conditions* exemplified in Figure 2. Each game is paired with the most concise utterance that establishes unique reference, using a vocabulary of 11 words: those corresponding to the six colors and four shapes, plus the word *shape*.

The dataset’s synthetic nature allows for highly controlled interventions on inherent properties of the world (e.g., color salience, Exp. 3), which can be difficult if not impossible to simulate with human learners in the real world.

Models

We experiment with two different types of speakers, an *RSA speaker* and a *literal speaker with context*. The **RSA speaker** reasons explicitly about the listener to generate utterances that optimize for successful communication. The **literal speaker with context** learns to directly map contexts to referring expressions. Based on the utterance produced by a speaker, the **evaluation listener** chooses one of the referents based on literal semantics learned from non-contextual data.

Evaluation listener The listener’s goal is to pick out the speaker’s intended target referent among three potential referents based on the speaker’s utterance. To do so, the listener L_{Eval} constructs a probability distribution over the referents

that encodes L_{Eval} ’s belief about which is the intended target. More specifically, given utterance u and referents r_1, r_2, r_3 , L_{Eval} will select referent r_k with probability

$$L_{\text{Eval}}(r_k | u, r_1, r_2, r_3) \propto \exp(\mathcal{L}(u, r_k)), \quad (1)$$

where \mathcal{L} is a neural-based semantic function mapping utterances to referents that is learned from data.

Given an utterance u and a referent r , \mathcal{L} outputs a probabilistic semantic value $p \in [0, 1]$ that represents the listener’s judgment about how well u applies to r . The probabilistic semantic value is computed as follows: given an utterance u and a referent r ,

$$\mathcal{L}(u, r) = \sigma(f_L(r)^T g(u)), \quad (2)$$

where f_L is an image encoder that maps an image to a d -dimensional vector, g is an utterance encoder that maps an utterance to a d -dimensional vector, and $\sigma : \mathbb{R} \rightarrow (0, 1)$ is the sigmoid function. d is a hyperparameter of our choosing; we choose $d = 1024$. f_L is a convolutional neural network (CNN; LeCun and Bengio, 1995); g is a unidirectional Gated Recurrent Unit (GRU; Cho et al., 2014).

When training \mathcal{L} , three training examples are extracted from each reference game. The ground-truth referring expression is paired up with each of the three referents: the expression paired with the target forms a positive example, and the other two pairs form negative examples.

Both the literal speaker and the RSA speaker are evaluated on the performance of this trained evaluation listener.

RSA speaker The RSA speaker S_{RSA} chooses a referring expression by recursively reasoning about a simulated listener’s interpretation. The speaker does so by considering how each utterance would be interpreted by an internal listener model L_{RSA} , which shares the same formulation as L_{Eval} (Eq. 1). After considering each possible utterance, the speaker chooses the most effective and efficient utterance, i.e., the one with the highest *utility*. An utterance’s utility U is defined as a trade-off between maximizing the likelihood of the listener identifying the correct target r_t and minimizing the cost $C(u)$ of producing utterance u :

$$U(u|t, r_1, r_2, r_3) = \log L_{\text{RSA}}(r_t | u, r_1, r_2, r_3) - C(u). \quad (3)$$

We define $C(u) := \lambda|u|$ to be a linear penalty on the length of an utterance, where $|u|$ denotes the number of words u contains. We choose $\lambda = 0.01$.

Once S_{RSA} has computed the utility for each possible utterance, it produces the utterance that maximizes the utility:

$$S_{\text{RSA}}(t, r_1, r_2, r_3) = \arg \max_{u \in \mathcal{U}} U(u|t, r_1, r_2, r_3). \quad (4)$$

L_{RSA} also learns semantics from data, encapsulated in its own semantic function \mathcal{L} . The learned nature of the semantic function introduces the potential for variation. To improve the robustness of L_{RSA} , we follow Wang et al. (2021) and model \mathcal{L} as the mean of an ensemble of n semantic functions $\{\mathcal{L}^{(1)}, \dots, \mathcal{L}^{(n)}\}$. All models in the ensemble are identical in architecture to L_{Eval} . We choose $n = 9$.

Literal speaker with context The literal speaker with context S_{Lit} is an image-to-text model with an encoder–decoder architecture. The image encoder f_S produces a vector representation of the context, i.e., the three referents. The decoder ℓ takes as input the output of f_S along with the target index to generate an utterance. More specifically, for a given reference game with referents r_1, r_2, r_3 , and a target index $t \in \{1, 2, 3\}$, S_{Lit} obtains an embedding of the reference game, \vec{h} , by first encoding the three referents individually with f_S and concatenating the encodings with the target index t :

$$\vec{h}(t, r_1, r_2, r_3) = [f_S(r_1); f_S(r_2); f_S(r_3); t] \quad (5)$$

Then, ℓ receives \vec{h} as input and generates an utterance:

$$S_{\text{Lit}}(t, r_1, r_2, r_3) = \ell(\vec{h}(t, r_1, r_2, r_3)) \quad (6)$$

Here, f_S is an image encoder with the same architecture as f_L , and ℓ is a GRU trained using the autoregressive language modeling objective with teacher forcing, which uses cross-entropy loss to penalize deviations from the ground-truth.

Training and implementation We adapted the neural module implementations of White et al. (2020). Each model was trained for 100 epochs with a batch size of 32, and optimized using Adam (Kingma & Ba, 2015) with a learning rate of 0.001 for the literal speaker and 0.01 for the semantic functions. For the literal speaker, we used validation accuracy as the model selection criterion. For the semantic functions, we used validation loss as the model selection criterion.¹

Experiment 1

We first investigate the degree to which the literal and RSA speakers choose overinformative referring expressions when all colors are equally likely to occur with all objects. This imitates the distribution of non-color-diagnostic objects, for which color overmodification is common (Sedivy, 2003).

Setup We created a dataset of 75,000 reference games, in which all color–shape combinations appear at a uniform frequency. 55,000 ($\approx 73\%$) of these were used to train the neural agents, and the remaining were reserved for evaluation.

The reference games used for training were split into 11 subsets of equal size: one for training the literal speaker S_{Lit} , one for training the semantic function \mathcal{L} in L_{Eval} , and nine for training the ensemble of semantic functions in S_{RSA} ’s internal listener L_{RSA} . The games were generated from the same distribution, but no two models were trained on identical data.

Results We start by noting that both speaker models have successfully learned to communicate the target to the evaluation listener with high accuracy (S_{Lit} : 93%, S_{RSA} : 94%). With this, we can turn to inspect the overmodification strategies the speaker models employ to achieve this accuracy.

¹All code and analyses for this paper can be found at github.com/feifang24/overmod-from-pragmatic-learning.

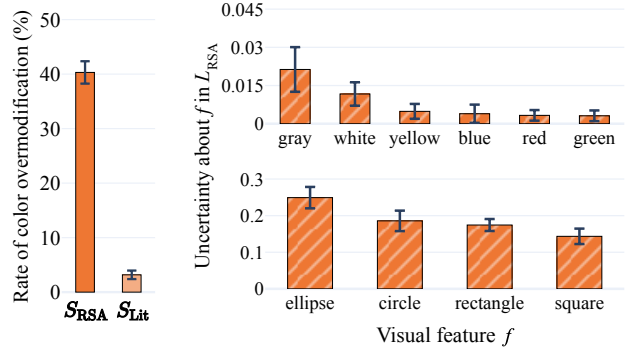


Figure 3: Results of Exp. 1 showing speakers’ rates of color overmodification (left), and the uncertainty associated with all visual features in the internal listener model L_{RSA} (right). Error bars correspond to 95% confidence intervals across 5 random initializations during training.

To investigate speakers’ color overmodification behavior, we focus on the rate of color in addition to shape mention (e.g., “blue square”) when color is neither sufficient nor necessary to identify the target (i.e., shape-needed condition in Figure 2). While strictly unnecessary, both speakers use color to overmodify, but the RSA speaker does so at a much higher rate (see Figure 3, left). The results are borne out as a significant main effect of the speaker on the rate of overmodification in a linear regression model ($\beta = .37$, $\text{SE} = .01$, $p < .001$). This suggests that overmodification with color is specifically encouraged in speakers that explicitly reason about listeners.²

While the literal speaker remains a black box, we can inspect the RSA speaker’s internal listeners to investigate what gives rise to the rate of color overmodification. More specifically, we can look at which visual features of the target are most or least likely to prompt the speaker to overmodify, and relate these patterns to the listener’s understanding of these features. We define *uncertainty* in the listener’s interpretation as the extent to which the outputs of the listener’s probabilistic semantic function \mathcal{L} deviate from the outputs from a truth-conditional semantics, which are 1 when the utterance applies to a referent and 0 otherwise.³ Figure 3 (right) shows substantial variation in feature uncertainty, despite balanced training data—a finding we return to in Exp. 3.

In alignment with prior work, we find that overmodification rates are highest where there is high uncertainty for shape and low uncertainty for color (Degen et al., 2020). S_{RSA} overmodifies with color most frequently when the target is an ellipse, which is associated with the highest uncertainty among all shapes (Figure 3 bottom right). S_{RSA} overmodifies with color least frequently when the target is gray, which is asso-

²Future work should explore how the rate of overmodification in these models is further affected by increased scene variation (e.g., Koolen et al., 2011).

³Treating deviation from end points as uncertainty is consistent with the probabilistic interpretation of real-valued semantics in (Degen et al., 2020), as well as with the classifier training objective for the literal listeners.

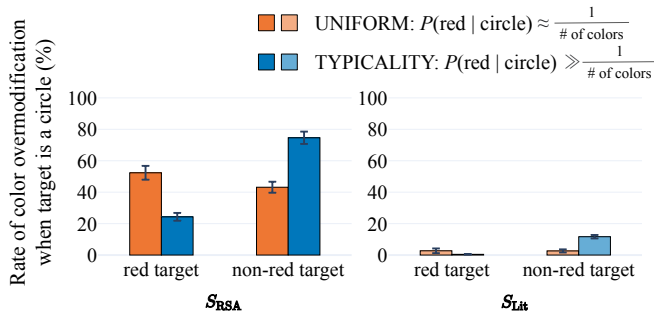


Figure 4: Rates of color overmodification when referring to red and non-red circles, after learning from uniformly colored circles (Exp. 1) and circles that are typically red (Exp. 2).

ciated with the highest uncertainty among all colors (Figure 3 top right). This observation is supported by a significant interaction of color and shape uncertainty on the rate of overmodification using a linear regression model ($\beta = -24.47$, $SE = 10.26$, $p < .05$)

Finally, we observe that when both color and shape are necessary for unique identification of the target (*both needed* condition, Figure 2), S_{RSA} and S_{Lit} achieve comparable listener accuracy ($90.43 \pm 0.20\%$ and $90.31 \pm 0.56\%$, respectively). However, when color is strictly not necessary (*shape needed*) but S_{RSA} overmodifies to a high degree, S_{RSA} achieves a higher listener accuracy than S_{Lit} ($87.56 \pm 0.34\%$ and $85.32 \pm 0.97\%$, respectively), borne out by a significant interaction between speaker and context condition ($\beta = .02$, $SE = .01$, $p < .01$). This suggests that overmodification might be an effective strategy that the RSA speaker leverages to improve communicative accuracy.

Experiment 2

Exp. 1 establishes that both the RSA speaker and the literal speaker generate referring expressions that overmodify with color, although at different rates. The speaker models were trained on underlying data where each color was uniformly paired with each shape, which is a commonly used feature distribution for models trained on synthetic data (e.g., ShapeWorld (Kuhnle & Copestake, 2017) or CLEVR (Johnson et al., 2017)). Conceptually, the uniform distribution of features makes color non-diagnostic for all object types (Tanaka & Presnell, 1999). In our second experiment, we investigate how the speaker models’ overmodification changes when certain objects are more likely to occur in some colors than in others. This setup imitates the distribution of color-diagnostics in the real world; humans refer to such objects overinformatively when they are atypically-colored but less so when typically colored (Westerbeek et al., 2015; Sedivy, 2003; Degen et al., 2020; Kreiss & Degen, 2020).

Setup We repeat the experimental protocol from Exp. 1 with a new dataset, TYPICALITY, in which the frequency of certain color–shape combinations deviates from uniform.

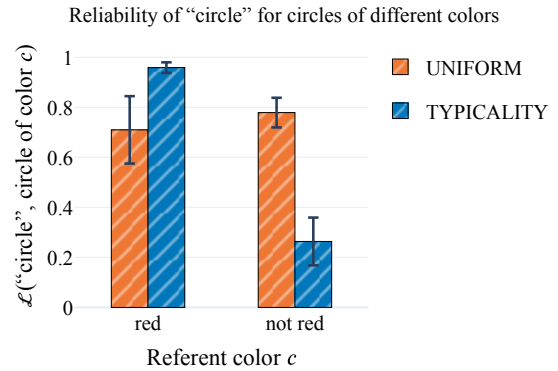


Figure 5: The internal listener model L_{RSA} ’s judgment about how well the utterance “circle” applies to circles of different colors, after learning from UNIFORM and TYPICALITY.

Specifically, 90% of all target circles are red, making circles in any other color atypical. We then examine how the overmodification behavior of a speaker trained on TYPICALITY differs from that of a speaker trained on the dataset from Exp. 1, which will constitute the UNIFORM condition.

Results Figure 4 shows how the data distribution impacts each speaker’s frequency of color overmodification when referring to a circle. Qualitatively, both speakers’ rates of overmodification decrease as the color becomes more typical (red in this case) and vice versa, resembling the empirical findings from prior literature (e.g., Westerbeek et al., 2015). However, as in Exp. 1, S_{RSA} overmodifies significantly more frequently than S_{Lit} , regardless of object typicality. These observations are borne out in a linear regression analysis as main effects of speaker (S_{RSA} vs. S_{Lit} ; $\beta = .63$, $SE = .02$, $p < .001$) and color of the circle (red vs. non-red; $\beta = -.11$, $SE = .02$, $p < .001$), as well as a significant interaction ($\beta = -.39$, $SE = .02$, $p < .001$).

Considering the range of color overmodification rates in the empirical typicality literature (e.g., typical: $\approx 20\%$; atypical: $\approx 75\%$ in Westerbeek et al. (2015), and typical: $\approx 25\%$; atypical: $\approx 55\%$ in Degen et al. (2020)), the color production rates of S_{RSA} are also quantitatively expected, whereas S_{Lit} underpredicts those rates.

We now again turn to investigate what gives rise to the overmodification behavior in S_{RSA} by examining its internal listeners’ semantic functions, which represent the listeners’ judgment about how well an utterance applies to a referent. When a listener is trained on TYPICALITY, the listener believes that “circle” applies well to a circle when it is of a typical color (i.e., red) and poorly otherwise (Figure 5). In contrast, when the listener is trained on the UNIFORM data, the listener perceives circles of all colors similarly. In other words, the listener’s semantic function significantly changes with the color–shape distribution it is exposed to during learning. This is borne out in a linear model as a significant interaction between dataset and red vs. non-red circles ($\beta = -.76$,




Figure 6: Example reference games from the HIGH-SALIENCE (Exp. 1) and LOW-SALIENCE conditions, respectively, both with the ground-truth utterance “red circle.”

$SE = .03, p < .001$). The results suggest that the overmodification asymmetry observed as a typicality effect can be captured simply by exposing a learning pragmatic agent with a non-uniform color distribution.

The results of the listener investigation are in line with the assumed semantics in prior work where the unmodified utterance (e.g., “banana”) is mapped with higher certainty to typically colored objects (yellow bananas) than to atypically colored ones (blue bananas) (Degen et al., 2020).

In summary, Exp. 2 shows that properties of the underlying data distribution play an important role in a speaker’s patterns of overmodification. While the overall rate in S_{RSA} is more compatible with previously observed human data, both speakers exhibit the same qualitative pattern of overmodification.

Experiment 3

The previous experiments explored whether color overmodification patterns from prior literature naturally emerge from neural speakers learning in certain environments where feature frequencies vary. The structured manipulation of the learning data led to changes in the learned uncertainties over features, which gave rise to the expected overmodification patterns. However, variation in uncertainty cannot simply be reduced to differences in frequency (e.g., see the uncertainty variation between *gray* and *red* in Figure 3). We now explore how overmodification changes when we leave the overall data frequency constant but make color more difficult to identify, rendering it less salient⁴. We thus generated a novel hypothetical world where an object’s color is determined by only a single colored pixel (right context in Figure 6). Following prior work, we expect that the reduced salience should reduce overinformative color use (Davies & Katsos, 2013; Sedivy, 2003; Rubio-Fernández, 2016; van Gompel et al., 2019).

Setup We use the dataset from Exp. 1 as the HIGH-SALIENCE condition and construct a novel variant of it, the LOW-SALIENCE condition (see Figure 6). While a shape’s color in the high-salience condition can be determined by randomly sampling *any* pixel within the shape, color in the low-salience condition needs to be extracted from one specific single pixel that is randomly placed within the shape.

⁴Salience is generally underspecified. Here, we take a simplistic approach and operationalize salience as ease of feature extraction.

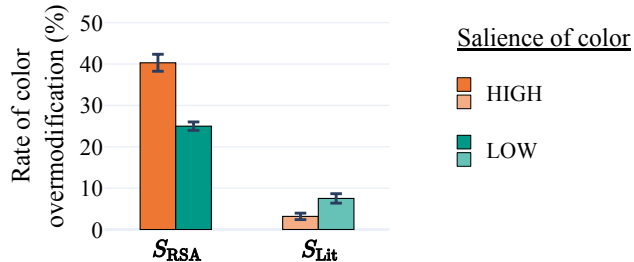


Figure 7: Frequency of color overmodification in the two speakers, trained on HIGH-SALIENCE (Exp. 1) and LOW-SALIENCE (Exp. 3), respectively.

Results As seen across experiments, S_{RSA} overmodifies with color more frequently than S_{Lit} regardless of the salience of color in the data distribution they learn from (Figure 7). However, in this experiment, S_{RSA} and S_{Lit} even come apart qualitatively. As color becomes less salient, S_{RSA} ’s frequency of color overmodification drops, while S_{Lit} ’s increases. These observations are borne out in a linear regression analysis with main effects of speaker ($\beta = .37, SE = .01, p < .001$), salience ($\beta = .04, SE = .01, p < .001$), as well as their interaction ($\beta = -.20, SE = .01, p < .001$).

The behavior of S_{RSA} is consistent with prior work suggesting that speakers might especially overmodify with color because of its salient nature (van Gompel et al., 2019; Rubio-Fernández, 2016; Davies & Katsos, 2013; Sedivy, 2003). While S_{RSA} ’s overmodification patterns are aligned with prior work, S_{Lit} ’s are at odds with it.

Upon inspecting the S_{RSA} ’s internal listener interpretations, we see that color modifiers are associated with about 10 times more uncertainty when the listener learns from an environment with low color salience (0.094 ± 0.063) than from one with high color salience (0.008 ± 0.002). This difference is borne out as a main effect in a linear model ($\beta = .09, SE = .01, p < .001$). The results are consistent with the findings from Exp. 1, which show that the RSA speaker overmodifies infrequently when the redundant information is associated with high uncertainty (Degen et al., 2020).

Conclusion

When referring to objects in the world, speakers often choose to mention an object’s color even when it is not necessary for uniquely picking it out among competitors (e.g., Pechmann, 1989). We have shown that a range of these attested overmodification patterns emerge naturally in neural speaker agents that learn semantics from data and reason pragmatically about internal listeners. Our findings align with color overmodification patterns reported for human speakers, and help to deepen existing modeling results concerning these patterns. More generally, we hope to have shown that neural networks are powerful tools for exploring how learning in different environments can shape pragmatic language use.

References

- Andreas, J., & Klein, D. (2016). Reasoning about pragmatics with neural listeners and speakers. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* (pp. 1173–1182). Austin, Texas: Association for Computational Linguistics. doi: 10.18653/v1/D16-1125
- Arnold, J. E. (2008). Reference production: Production-internal and addressee-oriented processes. *Language and cognitive processes*, 23(4), 495–527.
- Brown-Schmidt, S., & Konopka, A. E. (2011). Experimental approaches to referential domains and the on-line processing of referring expressions in unscripted conversation. *Information*, 2(2), 302–326.
- Cho, K., van Merriënboer, B., Gulçehre, Ç., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN Encoder-Decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 1724–1734).
- Cohn-Gordon, R., Goodman, N. D., & Potts, C. (2018). Pragmatically informative image captioning with character-level inference. In *Proceedings of the 2018 conference of the North American chapter of the Association for Computational Linguistics: Human language technologies* (pp. 439–443). Stroudsburg, PA: Association for Computational Linguistics.
- Cohn-Gordon, R., Goodman, N. D., & Potts, C. (2019). An incremental iterated response model of pragmatics. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2019* (pp. 81–90).
- Dale, R., & Reiter, E. (1995). Computational interpretations of the Gricean maxims in the generation of referring expressions. *Cognitive Science*, 19(2), 233–263.
- Davies, C., & Katsos, N. (2013). Are speakers and listeners ‘only moderately Gricean’? An empirical response to Engelhardt et al. (2006). *Journal of Pragmatics*, 49(1), 78–106.
- Degen, J., Hawkins, R. D., Graf, C., Kreiss, E., & Goodman, N. D. (2020). When redundancy is useful: A Bayesian approach to “overinformative” referring expressions. *Psychological Review*, 127(4), 591.
- Engelhardt, P. E., Bailey, K. G., & Ferreira, F. (2006). Do speakers and listeners observe the Gricean Maxim of Quantity? *Journal of Memory and Language*, 54(4), 554–573.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998.
- Franke, M. (2009). *Signal to act: Game theory in pragmatics*. Institute for Logic, Language and Computation Amsterdam.
- Fried, D., Andreas, J., & Klein, D. (2018). Unified pragmatic models for generating and following instructions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)* (pp. 1951–1963). New Orleans, Louisiana: Association for Computational Linguistics.
- Gatt, A., van Gompel, R. P., Krahmer, E., & van Deemter, K. (2011). Non-deterministic attribute selection in reference production. In *Proceedings of the CogSci workshop Production of Referring Expressions*.
- Golland, D., Liang, P., & Klein, D. (2010). A game-theoretic approach to generating spatial descriptions. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing* (pp. 410–419).
- Goodman, N. D., & Frank, M. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20.
- Goodman, N. D., & Stuhlmüller, A. (2013). Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science*, 5(1), 173–184.
- Hawkins, R., & Goodman, N. (2017). Why do you ask? The informational dynamics of questions and answers. *PsyArXiv*.
- Heller, D. (2020). The production and comprehension of referring expressions: Definite description. *Language and Linguistics Compass*, 14(5), e12370.
- Jager, G. (2007). Game dynamics connects semantics and pragmatics. *Game Theory and Linguistic Meaning*, 89–102.
- Johnson, J., Hariharan, B., Van Der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., & Girshick, R. (2017). Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 2901–2910).
- Kao, J. T., Wu, J. Y., Bergen, L., & Goodman, N. D. (2014). Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences*, 111(33), 12002–12007.
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30(1), 1–45. doi: 10.1007/s10988-006-9008-0
- Kingma, D. P., & Ba, J. (2015). Adam: A method for Stochastic Optimization. In Y. Bengio & Y. LeCun (Eds.), *3rd International Conference on Learning Representations*.
- Koolen, R., Gatt, A., Goudbeek, M., & Krahmer, E. (2011). Factors causing overspecification in definite descriptions. *Journal of Pragmatics*, 43(13), 3231–3250.
- Kreiss, E., & Degen, J. (2020). Production expectations modulate contrastive inference. In *Proceedings of the Annual Meeting of the Cognitive Science Society*.
- Kuhnle, A., & Copestake, A. A. (2017). Shapeworld - A new test methodology for multimodal language understanding. *CoRR, abs/1704.04517*.
- LeCun, Y., & Bengio, Y. (1995). Convolutional networks for images, speech, and time series. In M. A. Arbib (Ed.), (p. 255–258).

- Lewis, D. K. (1969). *Convention: A philosophical study*. Wiley-Blackwell.
- Long, M., Moore, I., Mollica, F., & Rubio-Fernandez, P. (2021). Contrast perception as a visual heuristic in the formulation of referential expressions. *Cognition*, *217*, 1048–79.
- Mitchell, M., Reiter, E., & Van Deemter, K. (2013). Typicality and object reference. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 35).
- Monroe, W., Hawkins, R. X., Goodman, N. D., & Potts, C. (2017). Colors in context: A pragmatic neural model for grounded language understanding. *Transactions of the Association for Computational Linguistics*, *5*, 325–338.
- Monroe, W., Hu, J., Jong, A., & Potts, C. (2018). Generating Bilingual Pragmatic Color References. In *Proceedings of the 2018 Conference of the North American chapter of the Association for Computational Linguistics: Human language technologies* (pp. 2155–2165). Stroudsburg, PA: Association for Computational Linguistics.
- Nadig, A. S., & Sedivy, J. C. (2002). Evidence of perspective-taking constraints in children's on-line reference resolution. *Psychological Science*, *13*(4), 329–336.
- Nie, A., Cohn-Gordon, R., & Potts, C. (2020). Pragmatic Issue-Sensitive Image Captioning. In *Findings of the Association for Computational Linguistics: EMNLP 2020* (pp. 1924–1938). Online: Association for Computational Linguistics.
- Pechmann, T. (1989). Incremental speech production and referential overspecification. *Linguistics*, *27*(1), 89–110.
- Rubio-Fernández, P. (2016). How redundant are redundant color adjectives? An efficiency-based analysis of color overspecification. *Frontiers in Psychology*, *7*, 153.
- Rubio-Fernandez, P. (2021). Color discriminability makes over-specification efficient: Theoretical analysis and empirical evidence. *Humanities and Social Sciences Communications*, *8*(1), 1–15.
- Rubio-Fernandez, P., Terrasa, H. A., Shukla, V., & Jara-Ettinger, J. (2019). Contrastive inferences are sensitive to informativity expectations, adjective semantics and visual salience. *PsyArXiv*.
- Sedivy, J. C. (2003). Pragmatic versus form-based accounts of referential contrast: Evidence for effects of informativity expectations. *Journal of Psycholinguistic Research*, *32*(1), 3–23.
- Sedivy, J. C. (2005). 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* (p. 345). MIT press.
- Syrett, K., Kennedy, C., & Lidz, J. (2009). Meaning and Context in Children's Understanding of Gradable Adjectives. *Journal of Semantics*, *27*(1), 1–35.
- Tanaka, J. W., & Presnell, L. M. (1999). Color diagnosticity in object recognition. *Perception & Psychophysics*, *61*, 1140–1153.
- van Gompel, R. P., van Deemter, K., Gatt, A., Snoeren, R., & Krahmer, E. J. (2019). Conceptualization in reference production: Probabilistic modeling and experimental testing. *Psychological Review*, *126*(3), 345.
- van Tiel, B., Franke, M., & Sauerland, U. (2021). Probabilistic pragmatics explains gradience and focality in natural language quantification. *Proceedings of the National Academy of Sciences*, *118*(9).
- Vedantam, R., Bengio, S., Murphy, K., Parikh, D., & Chechik, G. (2017). Context-aware captions from context-agnostic supervision. In *CVPR*.
- Waldon, B., & Degen, J. (2021). Modeling cross-linguistic production of referring expressions. In *Proceedings of the Society for Computation in Linguistics* (Vol. 4).
- Wang, R., White, J., Mu, J., & Goodman, N. (2021). Calibrate your listeners! Robust communication-based training for pragmatic speakers. In *Findings of the Association for Computational Linguistics: EMNLP 2021* (pp. 977–984).
- 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.
- White, J., Mu, J., & Goodman, N. D. (2020). Learning to refer informatively by amortizing pragmatic reasoning. *arXiv preprint arXiv:2006.00418*.