

Semantic Context Effects on Color Categorization

Rony Kubat (kubat@media.mit.edu)

MIT Media Lab, 20 Ames Street

Cambridge, MA 02139 USA

and **Daniel Mirman (mirmand@einstein.edu)**

Moss Rehabilitation Research Institute, Albert Einstein Healthcare Network

1200 W. Tabor Rd.

Philadelphia, PA 19141

and **Deb Roy (dkroy@media.mit.edu)**

MIT Media Lab, 20 Ames Street

Cambridge, MA 02139 USA

Abstract

A number of recent theories of semantic representation propose two-way interaction of semantic and perceptual information. These theories are supported by a growing body of experiments that show widespread interactivity of semantic and perceptual levels of processing. In the current experiment, participants classified ambiguous colors into one of two color categories. The ambiguous colors were presented either as a color patch (no semantic context), as an icon representing an object that is strongly associated with one of the color options, or as a word referring to such an object. Although the iconic and lexical contexts were incidental and irrelevant to the color categorization task, participants' responses were consistently biased toward the context color. These results extend previous findings by showing that lexical contexts, as well as iconic contexts influence color categorization.

Keywords: Color perception; Semantics; Interactive processing;

A number of current theories of the representation of word meanings propose that these representations are composed of perceptual and motor representations relevant for the word meaning. Such theories may be called *grounded* because they “ground” concepts in sensory/perceptual-motor/action representations (Barsalou, Simmons, Barbey, & Wilson, 2003; Roy, 2005); they may also be called *distributed* (McClelland & Rogers, 2003) because they propose that semantic knowledge is distributed throughout the perceptual-motor systems. Under such views, semantic knowledge of, for example, the meaning of color names involves precisely the same brain regions and representations as perception of color. In other words, the mechanisms of color perception are an intrinsic part of semantic knowledge. Conversely, this view also suggests that semantic knowledge is an intrinsic part of color perception.

This latter point is consistent with interactive (McClelland, 1993) and Bayesian (Geisler & Diehl, 2003) theories of perception and cognition, which propose that top-down information (such as semantic knowledge) directly influences lower-level perceptual processes. Visual perception, in particular, is a highly interactive process. The perception of illusory contours is guided by top-down expectations (Lee & Nguyen, 2001) and recognition of objects is influenced by scene contexts (Bar, 2004). Similarly, color perception appears to be influenced by memory of an object's color: when participants

were asked to adjust the color of a fruit object until it appeared achromatic, their judgments were consistently shifted away from gray in a direction opposite to the typical color of the fruit (Hansen, Olkkonen, Sebastian, & Gegenfurtner, 2006). That is, the participants' perception of the color of a displayed fruit was biased by the conceptual knowledge of the typical color of that fruit.

It has long been known that color perception is not merely a matter of the wavelengths of light hitting the retina—if it were, changes in illumination would radically alter the perceived colors in the environment. There is now direct evidence that top-down knowledge is used to improve color constancy (Mitterer & de Ruiter, 2008), much like context and word knowledge can be used to adjust speech sound categories (Norris, McClelland, & Cutler, 2003; Mirman, McClelland, & Holt, 2006). The use of context-dependent expectations gives the perceptual system robustness in the face of changing environments, which is a central principle of interactive (McClelland, 1993) and Bayesian (Geisler & Diehl, 2003) theories.

The paradigm for studying the effects of conceptual knowledge on color perception presented here is closely related to the Stroop (1935) paradigm. In the Stroop paradigm, two sources of information (traditionally, printed word and ink color) may be congruent or conflicting, and the interference caused by conflicting information reveals the interactions between processing of these different sources of information. In other words, the traditional Stroop paradigm tests the interaction of conceptual information from lexical context and perceptual information from color perception. The present experiment uses a Stroop-like paradigm to test the interaction of conceptual information (i.e., iconic or lexical contexts) with color information. Thus, the critical principles of information integration that are at work in the Stroop paradigm are the same as those at work in the present experiment.

We extend previous studies of the effect of context on color perception by testing whether in addition to iconic context, lexical context influences perception of ambiguous colors (i.e., ones that fall half-way between two basic color categories (Berlin & Kay, 1969)). If color perception is part of an interactive and/or a single integrated perceptual-cognitive system, simple color categorization will be influenced by high

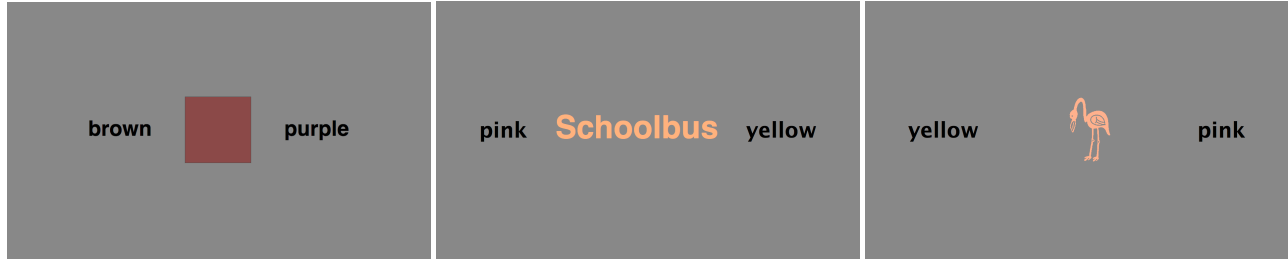


Figure 1: Screen captures from the experiment. From left to right: control, lexical context, iconic context. Colors in this figure are not representative of the experimental stimulus due to variations in the display or printing of the figure.

level semantic contexts. In effect, an ambiguous brown-purple patch will *appear* more brown when in the semantic context of “chocolate”, and the same color will *appear* more purple in the context of “eggplant.” By using iconic and lexical contexts, this experiment engages abstract, high-level semantic representations that are significantly removed from the low-level perceptual processes involved in color perception. Although both iconic and lexical contexts activate semantic representations, how those representations are activated is somewhat different: iconic contexts activate representations through object recognition, more directly than lexical contexts, which activate through language. Testing both iconic and lexical contexts within the same experiment makes the results directly comparable. An influence of these semantic contexts on color classification would demonstrate comprehensive interactivity of the perceptual-semantic system for the color domain.

Experiment

Across eight different pairs of basic color categories, English-speaking subjects chose which of two basic color terms better described an ambiguous color. Two different semantic context variations were tested. In the lexical-context condition a word with strong single color association (for example, the word “chocolate” for brown, and “eggplant” for purple) was colored in the ambiguous color. Similarly, in the iconic context condition a high contrast icon of an object commonly associated with a particular color category was presented in the ambiguous color. For the control condition, the ambiguous color was presented as a square swatch. If semantic context influences color perception, then participants should be more likely to label the ambiguous color consistent with the lexical or iconic context. That is, an ambiguous color between brown and purple (for example) should be labeled “brown” more often in the icon or word context of “chocolate” and labeled “purple” in the icon or word context of “eggplant”. Figure 1 shows example displays for each condition. A more detailed exposition of this experiment can be found in (Kubat, 2008).

Ambiguous Color Calibration

A calibration pilot study was conducted to identify the ambiguous colors. For each color pair in Table 1, a rectangular swatch was presented flanked by the two basic color terms.

Below the swatch and color labels, a slider allowed participants to change the mix between the two colors that filled the center swatch. Participants were instructed to use the slider to find the color that is perfectly ambiguous between the two color terms.

CIELab values for the focal colors i and j were chosen from the World Color Survey (Kay, Berlin, Maffi, & Merrifield, 2003) stimulus palette, clamped to the sRGB color space. The CIELab value of the center (ambiguous) color swatch was determined as:

$$L^* = \frac{iL^* + jL^*}{2}, \quad a^* = (1 - \alpha)i_{a^*} + \alpha j_{a^*}, \quad b^* = (1 - \alpha)i_{b^*} + \alpha j_{b^*}$$

where α is the slider value, ranging over $[0, 1]$. The luminance of the ambiguous color was fixed in order to minimize biasing based on perceived brightness and minimize perceptual contrast effects due to the experimental stimuli being presented against a neutral gray background.

Five ambiguous colors were later generated for each color pair using the alpha values of $\mu_{ij} + \beta\sigma_{ij}$, with $\beta \in \{-1.5, -0.75, 0, 0.75, 1.5\}$, where μ_{ij} and σ_{ij} are the mean and standard deviation of slider values chosen by the participants for color pair $\langle i, j \rangle$. By choosing these five deviations, the experiment control case is forced toward 50%, making any context effects discernible despite individual variations in color category boundaries.

Thirteen participants from the MIT community completed the calibration pilot study.

Participants

Twenty-three participants from the MIT community completed the main experiment (mean age = 37.5; 8 male, 15 female). All subjects were proficient speakers of English and reported normal color vision. None had participated in the calibration pilot study.

Materials and Procedure

The experiment was performed in a windowless, dimly lit room illuminated at approximately 3200°K. Participants were given time to adjust to the ambient lighting in the room before any color-related tasks were performed. Stimuli were presented with custom-written software on an Apple Macintosh computer and data recorded to a relational database. A

| Color Pair (<i>A</i> , <i>B</i>) | | Semantic Context | | | |
|---------------------------------------|--------|--------------------------|--------------------------|-------------------------|-------------------------|
| | | <i>A</i> lexical context | <i>B</i> lexical context | <i>A</i> iconic context | <i>B</i> iconic context |
| Red | Green | “Cherry” | “Broccoli” | Cherry | Leaf |
| Green | Blue | “Broccoli” | “Ocean” | Leaf | Waves |
| Blue | Red | “Ocean” | “Cherry” | Waves | Cherry |
| Pink | Yellow | “Flamingo” | “Schoolbus” | Flamingo | Corn |
| Yellow | Orange | “Schoolbus” | “Carrot” | Corn | Traffic Cone |
| Orange | Purple | “Carrot” | “Eggplant” | Traffic Cone | Grapes |
| Purple | Brown | “Eggplant” | “Chocolate” | Grapes | Log |
| Brown | Pink | “Chocolate” | “Flamingo” | Log | Flamingo |

Table 1: Color pairs used in the experiment and the lexical and iconic contexts used. Images of the icons used are shown in Figure 2.



Figure 2: Icons used in the experiment.

30-inch Apple Cinema Display was used as the display device. The monitor was calibrated to the sRGB standard (D65, 2.2 Gamma) using a ColorVision Spyder2Pro hardware color calibrator. All color stimuli were presented against a neutral gray background. The equipment for the ambiguous color calibration was identical to the experimental equipment.

In the main experiment, participants were presented with a color stimulus in the center of the screen, and two color labels in black text on the left and right. In order to proceed to the next screen, each participant chose which of the two labels better represented the center stimulus by pressing a key on the keyboard. Participants were instructed to proceed as quickly as they believed they could make an effective decision. To minimize color saturation effects, the stimulus color was only visible for 1500 milliseconds and was preceded with one second of neutral gray. Response time was recorded.

Ten trials were conducted for each condition (*A* lexical context, *B* lexical context, *A* iconic context, *B* iconic context, control) of each color pair (*A*, *B*) in Table 1. The ten trials presented the five ambiguous colors for each color pair in both color label orientations (left–right: *A*–*B* and *B*–*A*). Trials across all color pairs were divided into three blocks (control, lexical context, and iconic context) and randomized within block.

For the lexical context task, eight nouns were chosen based on their strong associations with a basic color category. Similarly, eight icons were selected to provide context—each with a strong color association. Table 1 details the words and icons used; Figure 2 shows the icons. This brought the total number of samples for each color pair to ten for the context-free (control) case, and twenty otherwise. The total number of decisions per participant collected during the experiment was 400 (8 color pairs, 10 control trials, 20 icon-context trials, and 20 work-context trials for each color pair).

Results

The mean of participants’ responses across color pairs, for each type of context, are shown in Figure 3. If a particular color continuum was not perceived as ambiguous by a particular participant (i.e., same color category response for all 10 control trials), that continuum was excluded from analysis for that participant. In the control condition, where the ambiguous color was presented as a square swatch, participants were about equally likely to choose each of the two responses (*A* response: mean = 52.7%, SEM = 1.9%). This slight bias

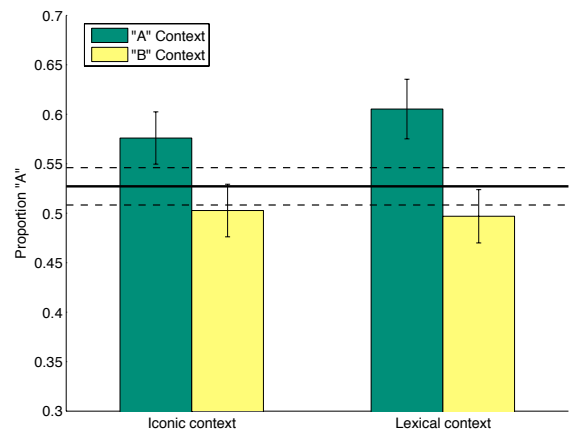


Figure 3: Mean “A” response for each type of context (error bars indicate 1SEM). Control (no context) condition response is illustrated by the solid line (dashed lines reflect 1SEM). Context biased color naming responses, causing an increase in context-consistent color naming for both lexical and iconic contexts.

| Color Pair (A, B) | | Percent Congruent | | n |
|----------------------|--------|-------------------|-----------------|----|
| | | Iconic context | Lexical context | |
| Red | Green | 33.42% | 37.37% | 19 |
| Green | Blue | 44.05% | 47.38% | 21 |
| Blue | Red | 58.18% | 60.00% | 22 |
| Pink | Yellow | 44.13% | 45.87% | 23 |
| Yellow | Orange | 62.83% | 64.35% | 23 |
| Orange | Purple | 63.53% | 67.94% | 17 |
| Purple | Brown | 60.87% | 60.87% | 23 |
| Brown | Pink | 63.61% | 62.78% | 18 |

Table 2: Percentage of participants who made a congruent choice (choosing *A* in an *A* context, or choosing *B* in a *B* context) separated by color pair. The number of included participants (those who found the control condition at least somewhat ambiguous) is *n*.

away from 50% reflects that though colors were very ambiguous, they were not perfectly ambiguous for every participant. Critically, the colors were sufficiently ambiguous to allow the predicted context effects to emerge. Iconic contexts had a reliable effect on responses ($t(22) = 2.90, p < 0.01$): when the ambiguous color was presented in the context of an *A*-biased icon participants were more likely to respond *A* (mean = 57.6%, SEM = 2.6%) and less likely to respond *A* in the context of a *B*-biased icon (mean = 50.2%, SEM = 2.7%). Lexical contexts also had a reliable effect on responses ($t(22) = 3.05, p < 0.01$). When the ambiguous color was presented in the context of an *A*-biased word, participants were more likely to respond *A* (mean = 60.5%, SEM = 3.0%) and less likely to respond *A* in the context of a *B*-biased word (mean = 49.7%, SEM = 2.7%). Table 2 separates the results by color pair. Five of the eight color continua showed strong context effects: greater than 50

Mean response times to the two context tasks were statistically identical to the control task indicating that there was no noteworthy additional cognitive loading added by the context. Response times separated between context-consistent and context-inconsistent responses (Figure 4) show a Stroop-like effect for iconic context (congruent: mean = 1539.51 ms, SEM = 82.43; incongruent: mean = 1632.92 ms, SEM = 93.91; $t(22) = 2.74, p < 0.05$), and an analogous though non-significant trend for lexical context (congruent: mean = 1640.87 ms, SEM = 99.16; incongruent: mean = 1680.72 ms, SEM = 95.36, $t(22) = 1.47, p = 0.155$). Response times for the control condition had mean 1600.54 ms, SEM = 91.20. Note that in a traditional Stroop paradigm congruency between color and lexical context are manipulated, whereas this analysis compares identical displays in which the participants produced congruent vs. incongruent responses. Incongruent responses were slower (statistically reliable only for iconic contexts), suggesting that they were produced in the face of competition from context information.

Discussion

These results demonstrate a semantic effect on color categorization. When ambiguous colors were presented in the con-

text of a word or iconic image, categorization of the colors was biased in the direction of the typical color of that pictured or named object. That is, participants’ semantic knowledge of the typical color of common objects shifted their perception of ambiguous colors to be closer to the typical color for that object. The knowledge that chocolate is typically brown shifted the perception of its color to be more brown. This result replicates previous findings of semantic effects on color perception (Hansen et al., 2006; Mitterer & de Ruiter, 2008), and extends them by showing that both iconic and lexical semantic contexts influence color categorization.

There are three interpretations of the observed phenomena: (1) context effects are a decision-level, feed-forward effect, (2) they are a result of top-down interaction effects, or (3) they due to shared representation of color between perceptual and conceptual processing.

Under the assumption of separate cognitive components, one dedicated to color perception and another dedicated to conceptual processing, the flow of information distinguishes between strictly feed-forward and interactive views of visual processing. In the feed-forward view, information flow is unidirectional from perception to conceptual processing. In the present experiment participants categorized ambiguous colors by choosing which of two color labels was a better label for the stimulus color. This label-based color categorization task requires access to semantic knowledge about the colors denoted by the labels. As a result, the task can be re-cast as semantic categorization—similar to semantic categorization experiments where participants categorize objects as “animals” or “plants”; in this case participants chose between “yellow” and “pink” as semantic categories. Under this view, one could argue that top-down feedback is not necessary to account for the results since all of the critical processing took place at the semantic level. Necessarily, any contextual ef-

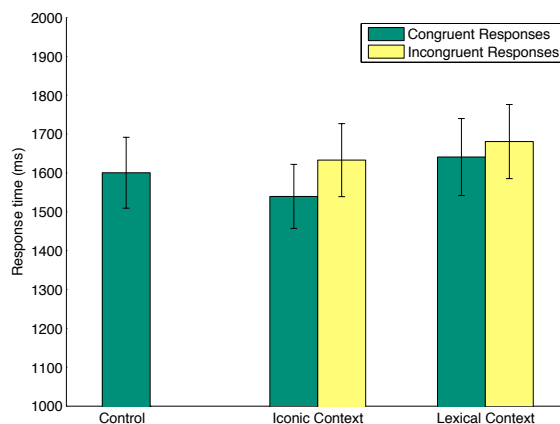


Figure 4: Mean response times for control, congruent and incongruent responses for both iconic and lexical contexts in forced choice tasks (error bars indicate 1SEM).

fects from icons or words take place at a post-perceptual decision level that integrates perceptual and conceptual information. This account is problematic with respect to recent results (Mitterer & de Ruiter, 2008) which show recalibration of color categories that could not be produced by decision-level integration of color and conceptual knowledge.

The interactive interpretation allows information to flow between components bidirectionally, with conceptual processing directly influencing perceptual processing; namely, by biasing the perception of color toward the context concept. The findings presented here are consistent with a long history of evidence of top-down effects in visual processing (Bar, 2004), as well as other domains (McClelland, 1993).

According to distributed and perceptually-grounded theories of semantic representation (Barsalou et al., 2003; McClelland & Rogers, 2003) semantic representations are based on the relevant perceptual representations. That is, the semantic representations of “yellow” and “pink” are, in large part, the perceptual representations of yellow and pink colors. The modules are not distinct entities, but rather intimately intertwined, sharing state. This view represents an extension of the interactive view in which bidirectional information flow blurs the distinction between components. Under the distributed/grounded view, conceptual effects on color perception are due to integration of conceptual and perceptual representations, rather than top-down feedback between distinct representational levels.

Whether one takes the interactive or integrated view, the present results add to the growing body of evidence that the perceptual-semantic system does not consist of isolated modular components. If just seeing the word *chocolate* or a simple icon representing it is enough to influence the perception of an ambiguous color, then the very notion of a strictly bottom-up color perception module that is independent of the perceiver’s knowledge and context is undermined. If there are processing components, they are richly interactive. Alternatively, it may be more appropriate to consider the perceptual-semantic system as an integrated whole with no truly independent components.

Acknowledgments

The authors would like to thank Matthew Ng and the members of Cogmac for their comments and suggestions. This paper is based upon work supported under a National Science Foundation Graduate Research Fellowship to RK and by National Institutes of Health grant F32HD052364 to DM.

References

Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, 5, 617-629.

Barsalou, L. W., Simmons, W. K., Barbey, A. K., & Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences*, 7, 84-91.

Berlin, B., & Kay, P. (1969). *Basic color terms: Their universality and evolution*. University of California Press.

Geisler, W. S., & Diehl, R. L. (2003). A bayesian approach to the evolution of perceptual and cognitive systems. *Cognitive Science*, 27, 379-402.

Hansen, T., Olkkonen, M., Sebastian, W., & Gegenfurtner, K. R. (2006). Memory modulates color appearance. *Nature Neuroscience*, 9, 1367-1368.

Kay, P., Berlin, B., Maffi, L., & Merrifield, W. R. (2003). *World color survey*. Center for the Study of Language and Information.

Kubat, R. (2008). *A context-sensitive meta-classifier for color-naming*. Unpublished master’s thesis, Massachusetts Institute of Technology.

Lee, T. S., & Nguyen, M. (2001). Dynamics of subjective contour formation in early visual cortex. *Proceedings of the National Academy of Sciences*, 98, 1907-1977.

McClelland, J. L. (1993). Toward a theory of information processing in graded, random, interactive networks. In D. E. Meyer & S. Kornblum (Eds.), *Attention & performance xiv: Synergies in experimental psychology, artificial intelligence and cognitive neuroscience* (p. 655-688). Cambridge, MA: MIT Press.

McClelland, J. L., & Rogers, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews Neuroscience*, 4, 310-322.

Mirman, D., McClelland, J. L., & Holt, L. L. (2006). An interactive hebbian account of lexically guided tuning of speech perception. *Psychonomic Bulletin & Review*, 13, 958-965.

Mitterer, H., & de Ruiter, J. P. (2008). Recalibrating color categories using world knowledge. *Psychological Science*, 19, 629-634.

Norris, D., McClelland, J. L., & Cutler, A. (2003). Perceptual learning in speech. *Cognitive Psychology*, 47, 204-238.

Roy, D. (2005). Grounding words in perception and action: computational insights. *Trends in Cognitive Science*, 9, 389-396.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18, 643-662.