
Is color an intrinsic property of object representation?

Galit Naor-Raz, Michael J Tarr[¶]§

Department of Psychology, (§ Department of Cognitive and Linguistic Sciences), Brown University, Box 1978, Providence, RI 02912, USA; e-mail: Michael_Tarr@brown.edu

Daniel Kersten

Department of Psychology, University of Minnesota, Minneapolis, MN 55455, USA
Received 3 June 2002, in revised form 28 October 2002; published online 6 June 2003

Abstract. The role of color in object representation was examined by using a variation of the Stroop paradigm in which observers named the displayed colors of objects or words. In experiment 1, colors of color-diagnostic objects were manipulated to be either typical or atypical of the object (eg a yellow banana versus a purple banana). A Stroop-like effect was obtained, with faster color-naming times for the typical as compared to the atypical condition. In experiment 2, naming colors on words specifying these same color-diagnostic objects reversed this pattern, with the typical condition producing longer response times than the atypical condition. In experiment 3, a blocked condition design that used the same words and colors as experiment 2 produced the standard Stroop-like facilitation for the typical condition. These results indicate that color is an intrinsic property of an object's representation at multiple levels. In experiment 4, we examined the specific level(s) at which color–shape associations arise by following the tasks used in experiments 1 and 2 with a lexical-decision task in which some items were conceptually related to items shown during color naming (eg banana/monkey). Priming for these associates was observed following color naming of words, but not pictures, providing further evidence that the color–shape associations responsible for the differing effects obtained in experiments 1 and 2 are due to the automatic activation of color–shape associations at different levels of representation.

1 Introduction

What physical properties of objects are encoded in their visual representations? It is uncontroversial that shape is critical to remembering and identifying things, but less clear what status surface properties such as texture or color have in the representation of objects. Answers to this latter question are not without disagreement—claims have been made for no significant effect (Biederman and Ju 1988) and for robust effects (Hayward and Williams 2000; Price and Humphreys 1989; Tanaka and Presnell 1999; Wurm et al 1993) of surface properties on visual recognition. Moreover, although most vision researchers would concede that color constitutes a component of internal object representations, almost every investigation into this question has contrasted the active *recognition*—naming, verification, recognition memory, or sequential-matching—of line drawings of objects with colored, shaded, or both colored and shaded images of objects. While such tasks may, appropriately, tap into the kinds of mechanisms used in everyday recognition, they are subject to an important caveat: object identification necessarily recruits both visual and nonvisual sources of information. Although information from multiple domains may affect any judgment about an object, recognition is more likely to draw on various sources of information in that many visual and semantic properties of an object are combined into a single entity (thing 'X'). Consequently, it may be difficult to establish the specific level or levels of processing (visual, conceptual, or lexical) at which the surface properties of an object affect its recognition (see Tanaka et al 2001).

[¶] Author to whom all correspondence and requests for reprints should be addressed.

Consider the canonical example of a yellow banana. At the visual level, the object ‘banana’ may be indexed by physiological measurements of the shape of bananas, eg the responses of shape-selective receptive fields in IT cortex (Tanaka et al 1991), quantitative measurements of the color yellow typical of bananas (Livingstone and Hubel 1984), and the flat, smooth texture of a banana’s skin. At the conceptual level, knowledge about bananas may include an amodal collection of features (Collins and Quillian 1969; Meyer and Schvaneveldt 1971) including descriptions of shape, eg “bent, elongated axis, tapered ends, and round cross-section”, the surface, eg “yellowish-green, brown speckles, flat, and smooth”, and function, eg “edible, sweet, eaten by monkeys”. At the lexical level, the lexical entry “banana” might be associated with the entry ‘yellow’ because the two are more likely to co-occur than “yellow” and “spinach” or “banana” and “purple”.

When one is asked to identify or otherwise recognize a banana, any and all of these sources of information are likely to come into play. For instance, generating the name “banana” in response to a colored picture of a banana may rely on matching visual properties of shape and color in the visual domain, identifying corresponding features in the conceptual domain, and activating the appropriate entry in the lexical domain. In the case where a specific color is strongly associated with a specific shape, object recognition may be based on both shape and color. In contrast, when color and shape are not closely associated, recognition must be based on shape alone. Thus, it is not surprising that many researchers have observed that generating the name of an object with a typical color, eg “banana”, is faster for colored pictures as compared to line drawings (Williams and Tanaka 2000). Likewise, it follows that when shape is less diagnostic (eg similarly shaped objects such as quadrupeds or fruit—Price and Humphreys 1989) or degraded (eg through occlusion or low-vision—Wurm et al 1993; Tanaka and Presnell 1999), diagnostic surface properties such as color (to the extent they are available) become more heavily weighted in recognition judgments. To make the point that different object properties are weighted appropriately to their diagnosticity, one could run a simple experiment—painting each of several differently shaped objects a very different color. No doubt, observers aware of the manipulation would use this color information and respond faster when the informative colors were present (Hayward and Williams 2000; Tarr and Bülthoff 1995).

What is *not* addressed in studies of recognition, however, is the degree to which an object’s shape and surface properties are automatically associated in its *representation*. In the thought experiment above, the identity of each object becomes statistically associated with a given color (in this case with a probability of 1.0)—whether this relationship is part and parcel of the object representation, particularly at the visual level, or is simply semantic knowledge about the object, is unknown. At issue is not whether we know that bananas are yellow (at some level) or even that the things that we call bananas have a particular surface appearance, but whether the visual representation of a banana *intrinsically* includes surface properties such as color. It is this last question about which researchers such as Biederman and Ju (1988), and Price and Humphreys (1989) disagree, yet fail to address directly in the naming and recognition paradigms used in their studies.

Here we attempt a different approach to this question. Using a variation of the Stroop paradigm (Stroop 1935), we ask observers to make judgments that *do not* directly invoke object identity. Consider a color image of a banana. If the task is to name the *color* of the banana, are responses faster when the banana is yellow as compared to when it is purple? A difference of this sort would be consistent with the classic finding that color naming responses are faster when colors appear on consistent color words as opposed to inconsistent color words (eg the word “red” in red ink versus the word “red” in green ink—Stroop 1935). More importantly, what would this result tell us about the visual representation of objects? In section 6 we will argue

that it suggests, at least for color-diagnostic objects, that surface properties such as color are visually represented concurrently with an object's shape.

In our 'object Stroop' experiment the participant need only name the color of the object (and indeed this color is the only non-black color in the display), yet there may be automatic processing of the object's shape (but not necessarily to the level of its identity). If so, *and* there is a specific color associated with the representation of this shape, that color may become active from the visual level on up and either facilitate or interfere with the generation of the appropriate color term. Note that this framework provides a more precise claim with regard to the status of color in the visual representation of objects. It hypothesizes that surface properties such as color are related *intrinsically* to shape properties and, therefore, cannot be processed independently. Thus, if an object's shape is accessed, there is concurrent activation of associated surface properties, apart from whether either shape or surface information is task-relevant (as in naming) or task-irrelevant (as in our experiment). It is whether this sort of processing actually occurs that we investigate.

2 Experiment 1

2.1 Methods

2.1.1 *Participants*. The participants in all four experiments consisted of Brown University undergraduate students who received either course credit or payment for their participation. In all four experiments, observers were counterbalanced for procedural purposes by using two groups. Experiment 1 included twenty-two participants.

2.1.2 *Stimuli*. In experiment 1 we used colored pictures showing everyday objects that were obtained from the Photodisc collection (Photodisc Inc., Seattle, WA) and modified in PhotoShop (Adobe Systems Inc., San Jose, CA). Seven basic, highly identifiable colors were used for all of the pictures: blue, red, green, yellow, orange, purple, and white. The objects were all *color-diagnostic* objects, defined as objects that tend to have at least one typical color strongly associated with them. For instance, a banana is strongly associated with the color yellow, while a pumpkin is strongly associated with the color orange. Altogether, thirty color-diagnostic objects were used (Appendix 1). Each of these objects was rendered in two different colors (both chosen from the seven colors listed above): one that was typical of the object and one that was an atypical color for the object. Thus, there was a total of sixty stimulus images. To ensure that our assignment of typical and atypical colors was accurate, we ran a pilot experiment in which we asked subjects to rate the typicality of the color shown on each object. Ratings for typical colors were uniformly high, while ratings for atypical colors were uniformly low. Thus, there was little ambiguity regarding the appropriateness of the colors used in each condition.

2.1.3 *Design and procedure*. In this and subsequent experiments all computer-run conditions were controlled by RSVP software (<http://www.tarrlab.org/RSVP/>) running on an Apple Macintosh computer. A CMU Button-Box equipped with a voice-trigger was used to record naming response times accurate to 1 ms at the onset of the spoken color name.

On each trial, a single image of an object was presented in the center of the screen. Participants were instructed to verbally name the *color* shown on the screen as quickly as possible, paying no attention to the object. Responses were hand-scored for correctness and response times were recorded with the voice trigger. The exact trial sequence was as follows: a fixation point appeared on the screen for 500 ms, followed by the stimulus, which remained on the screen until the participant responded. An inter-trial interval of 1000 ms occurred between trials. Participants were instructed to use only the seven basic colors (yellow, red, green, etc). The trials were counterbalanced across

two groups of participants, such that no single observer saw the same stimulus in both its typical and atypical forms.

2.2 Results and discussion

Across all conditions and participants there was a mean accuracy of 98.7% (8/630 were incorrect, 5 in the atypical condition and 3 in the typical condition). Incorrect responses as well as response times over 2500 ms were excluded from further analyses. A two-way ANOVA on response times with group and condition as factors revealed no main effect for group ($F < 1$), or interaction of group with condition ($F < 1$). Therefore, the two groups were collapsed in all subsequent analyses.

As illustrated in figure 1, mean response times were faster for objects that appeared in typical colors rather than atypical colors. A within-subjects one-way ANOVA confirmed this difference to be significant ($F_{1,20} = 10.77$, $p < 0.005$). The significant difference between the typical and atypical conditions indicates that color must be associated with an object's shape at some level of representation. More specifically, previous knowledge of the object's typical color interfered with the ability of observers to name the atypical color presented on the object.

What is not entirely clear is the form of this previous knowledge regarding color. In experiment 1, the stimulus was purely visual—an image of an object—although the response was necessarily verbal. Thus, generating a response required linking visual processes to conceptual and ultimately lexical processes. This chain of processing was only necessitated in the color domain—nothing in the task requires that participants process shape at the visual level or beyond. Yet, the color typicality effect we obtained indicates that shape was accessed insofar as colors associated with specific shapes affected performance. That is, *automatic* processing of the irrelevant shape occurred, and as such also led to activation of associated color information at some level of object representation. Indeed, it is often the case that when specific perceptual information is processed by default (as in face recognition), interference is obtained in a second task for which this information is not relevant (Gauthier et al 2003). The same principle applies here—automatic processing of a given object shape activated information about the typical color of that object, and in doing so facilitated or interfered with the color-naming task. The most plausible account is that this shape knowledge was most active at the visual level (because the task did not require knowledge about the object or lexical activation of the object name). At the same time, we acknowledge that the results of experiment 1 are not definite on this point. That is, if automatic processing of the object occurs beyond the visual level, then the origins of the effect of the color–shape association obtained in experiment 1 may be either semantic or lexical in nature. However, our hypothesis is that the obtained color–shape association is *visual* and independent of semantic or lexical knowledge regarding an object's appearance.

It is possible to reverse this argument, asking how would the color–shape association manifest itself if the level of access provided to the observer consisted of a *word* denoting the object rather than its *picture*? In this case, processing would commence at the lexical level and then progress to the conceptual and visual levels. Thus, any color–shape association would most likely occur at the lexical, and perhaps conceptual, level. Would we obtain a Stroop-like effect, and would it be as strong as the one obtained in the current experiment? If the effect is different than that obtained in experiment 1, it would add further weight to our hypothesis that the color–shape associations arising from naming colors on pictures of objects are not the result of conceptual or lexical associations.

3 Experiment 2

It is well established that Stroop-like effects are obtained with colors placed on words (Klein 1964; Ménard-Buteau and Cavanagh 1984). Such studies have typically relied on a blocked design in which participants read through lists of entirely color-typical or color-atypical color–word pairs and the dependent measure is the time taken to name every color in the complete list. Thus, it is impossible to compare these previous results to the effect we obtained in experiment 1. Therefore, in experiment 2 we adopted the same mixed condition designed as that used in experiment 1, but with colored word stimuli rather than colored pictures of objects.

3.1 Methods

3.1.1 *Participants.* Experiment 2 included twenty participants.

3.1.2 *Stimuli.* In experiment 2 we used 30 words (36 point Helvetica Bold) depicting the same objects as in experiment 1. The same colors and color combinations were used as in experiment 1.

3.1.3 *Design and procedure.* The same procedure was employed as in experiment 1.

3.2 Results and discussion

Analyses were the same as those applied in experiment 1. An accuracy assessment yielded an accuracy level of 99% (6 errors out of 630 trials, all in the typical condition). A two-way ANOVA on response times with group and condition as factors revealed no main effect for group ($F < 1$), or interaction of group with condition ($F < 1$). Therefore, the two groups were collapsed in all subsequent analyses.

Figure 1 presents the results of experiment 2. Somewhat surprisingly, color-naming times were *shorter* for atypical color–word pairs than for typical ones. A one-way within-subjects ANOVA indicated that this difference was significant ($F_{1,19} = 9.63, p < 0.01$).

The findings of experiments 1 and 2 point in two different directions, indicating an interesting dissociation in the associations between color and shape. First, these results reinforce the point that surface properties such as color are strongly associated with an object's shape at some level of representation. Otherwise, no difference would have been observed between the typical and atypical conditions. Second, the fact that the atypical condition induced shorter color-naming times relative to the typical condition suggests that the critical difference between experiments 1 and 2 is *how* color and shape information is associated given different presentation modalities (pictures versus words). In experiment 1, the stimuli were exclusively visual (in the form of pictures). In contrast, in the current experiment, the stimuli were orthographic (in the form of printed words). Pictures necessarily activate visual representations at the earliest stages of picture processing, while words minimally activate their corresponding entries in the 'mental dictionary' known as the lexicon. In contrast to visual object representations, lexical representations capture both the form (in terms of sound structure and/or orthography up to meaning) and the syntactic role of a word (eg Foss and Hakes 1978).

We contend that it is this difference in initial access that leads to opposite effects in the two experiments. Specifically, color–shape associations in visual representations produce facilitation for color-name access because each object is represented by both its shape *and* its color. In contrast, color–shape associations that arise during lexical access, that is when searching for a particular lexical item that has the appropriate structure and meaning, will produce competition between the lexical entries for the name of the object and the name of the color, thereby slowing color-name access. This latter hypothesis is supported by the fact that every color-name error in experiment 2 occurred in the typical color condition—the case where one would predict the greatest competition between the object name and the color name and, consequently, an increased likelihood of naming errors.

4 Experiment 3

The results of experiment 2 contrast with the results of Klein (1964). As in experiment 2, Klein used words specifying color-diagnostic objects (although he included many mass nouns, eg “blood” and “grass”, while we used only count nouns) in a Stroop paradigm, having participants name the ink color of each word. A given name could appear in either its typical or an atypical color. Using a blocked list design, he obtained the standard Stroop effect in which color names were generated faster for consistent color–object pairs.

The difference in results between our experiment 2 and Klein’s study may be due to our mixed condition design as compared to his blocked condition design. That is, in experiment 2 participants did not know on a given trial whether the color–object pair would be consistent or inconsistent. Moreover, we measured color-naming time for individual trials, whereas in Klein’s study a single summed color-naming time was measured for each condition (a single list).

To investigate whether these methodological differences were at the root of the opposite effects obtained by us and Klein, we employed the stimuli from experiment 2 in a blocked list design similar to that used by Klein (1964).

4.1 Methods

4.1.1 *Participants.* Experiment 3 included twenty-eight participants.

4.1.2 *Stimuli.* Stimuli consisted of the same set of colored words (object names) as that used in experiment 2. Rather than using computer-controlled presentation, each word was presented as a printed item in a vertically oriented list (one word per line) on white paper. Response times were recorded manually with a stopwatch.

4.1.3 *Design and procedure.* Participants received different instructions for each list. For all lists, they were instructed to perform the task as quickly as possible, but to correct any perceived mistakes, and then immediately proceed onto the next item. In addition to the two color–object lists, two practice lists were used to allow participants to familiarize themselves with the task. The four lists and the associated instructions were very similar to those used by Klein (1964). The lists were presented to participants as follows:

- (i) Black-ink condition. Thirty words depicting colors were printed in black. The task was to read aloud each color name in the order they appeared in the list.
- (ii) Asterisk condition. Thirty rows of asterisks were printed in seven different colors. The task was to name the color of each row of asterisks.
- (iii) Typical condition. Thirty names of objects were printed in their typical colors. The task was to ignore the semantic content of the words and name the color of each word in the order it appeared in the list.
- (iv) Atypical condition. Thirty names of the same objects as in the typical condition were printed in atypical colors. The task was to ignore the semantic content of the words and name the color of each word in the order it appeared in the list.

Each participant was provided with the appropriate instructions before each condition (for details, see Klein 1964). A single overall response time per list was recorded by the experimenter with a stopwatch. Participants were divided into two groups, group 1 receiving the typical condition before the atypical condition, and group 2 received them in the reverse order. The two practice conditions were always presented to participants prior to either the typical or atypical conditions.

4.2 Results and discussion

Across all conditions, color naming was 98.5% accurate, with a total of 49 mistakes in 3360 trials. Errors were distributed as follows: 8 in the typical condition, 14 in the atypical condition, 26 in the asterisk condition, and 1 in the black-ink condition. A two-way ANOVA with group and condition as factors revealed no main effect for

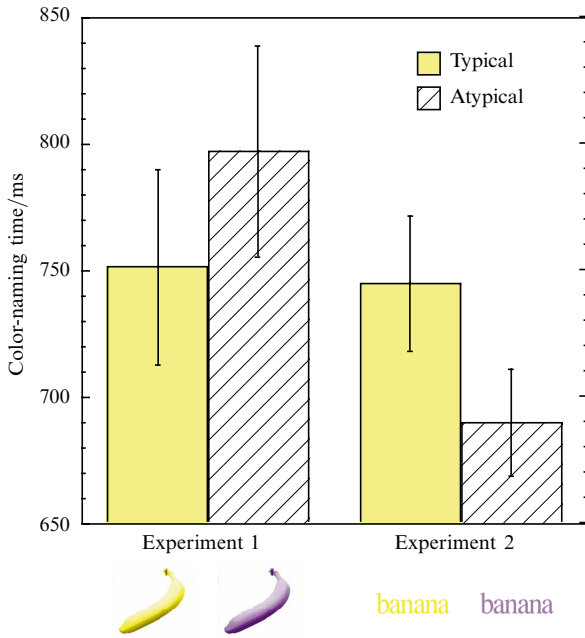


Figure 1. Results of experiments 1 (pictures) and 2 (words). Error bars show between-subjects standard error.

group ($F < 1$), or interaction of group with condition ($F < 1$). Therefore, the two groups were collapsed in all subsequent analyses.

A within-subjects ANOVA was performed with condition as the only factor. This analysis revealed a significant difference between conditions ($F_{3,27} = 60.58$, $p < 0.0001$), a pattern that is presented in figure 2. A separate ANOVA including only the typical and atypical conditions revealed a significant difference in this critical comparison ($F_{1,27} = 4.88$, $p < 0.05$). Thus, overall naming times in the atypical condition

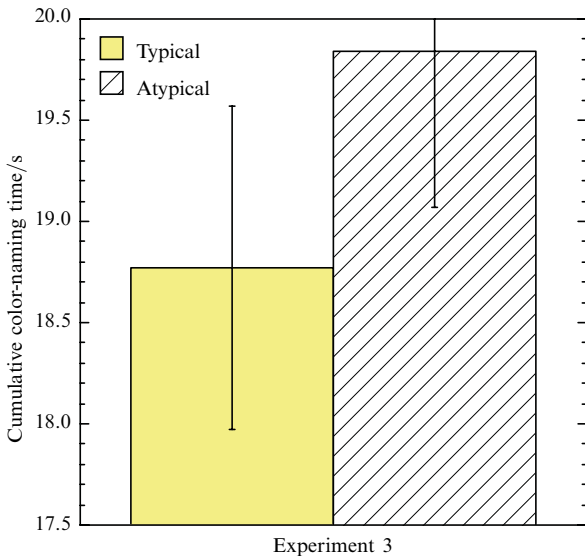


Figure 2. Results of experiment 3 (words in a blocked design). Error bars show between-subjects standard error.

were reliably different from overall naming times in the typical condition, a pattern consistent with Klein's (1964) results.

Critically, we were able to replicate Klein's study (1964) with the color-diagnostic object set and color–shape pairings used in experiments 1 and 2. Thus, we can be fairly certain that the differences between experiment 2 and Klein's original results lie in differences in methodology, not idiosyncratic properties of our stimuli (eg the nature of the color–shape associations). In some sense the difference between experiment 2 and Klein's study (as well as our experiment 3) is not surprising—under blocked conditions, the participant is certainly aware that each to-be-named color is either consistent or inconsistent within the entire list. Consequently, the participants can deploy attention in a somewhat different and more strategic manner than in a mixed condition design. Specifically, given only atypical colors on the words, participants are aware that the words provide no task-relevant information and therefore may actively attempt to inhibit lexical processing. In contrast, given only typical colors on the words, they are aware that the words provide additional information relevant to the task, and therefore may attempt to initiate lexical processing. Within a mixed condition design, such strategic biases provide no overall benefit to the participant and, furthermore, will cancel out over the course of the experiment.

5 Experiment 4

The core issue addressed by experiments 1 and 2 is what kind of information about an object is *automatically* activated during color naming of pictures or words. In experiment 1 we hypothesized that the perception of pictures of objects would automatically recruit visual object representations that included color–shape associations. The fact that we obtained a standard Stroop effect in experiment 1 supports this prediction, suggesting that color information associated with object shape was activated. However, the specific form of such associations cannot be isolated by this finding alone. To help isolate the different loci of color–shape associations in object representations, we ran experiment 2, in which we hypothesized that the perception of words specifying objects would automatically recruit lexical and, perhaps, conceptual representations of objects. We obtained a reversed Stroop effect as compared to experiment 1, indicating that although color and shape were associated in both experiments, different kinds of object representations were automatically activated as a consequence of perceiving a picture or word. Taken together, the results of experiments 1 and 2 point to color–shape associations that arise from different forms of object representation. On the basis of both plausibility and the different patterns of interference effects across the two experiments, we hypothesize that pictures automatically recruit visual representations of objects while words automatically recruit lexical and conceptual representations.

To further clarify the nature of these color–shape associations, experiment 4 followed the picture color-naming task from experiment 1 or the word color-naming task from experiment 2 with a lexical-decision task (word/non-word judgment). Each lexical item was related to a picture or word shown during color naming (eg “monkey”, where either a picture of a banana or the word banana was shown previously), unrelated to any picture or word shown previously, or was a non-word. Our reasoning was that, if the effects obtained in experiment 1 were due to the activation of conceptual representations, we should obtain significant priming from banana(picture) to monkey(word). However, if the effects obtained in experiment 1 were due to the activation of visual representations of objects, we predicted no significant priming.⁽¹⁾ At the same time, we reasoned that if color–shape associations at the conceptual and lexical levels were

⁽¹⁾On the assumption that bananas and monkeys are not strongly related at the visual level—an assumption that would be incorrect if a participant had often seen monkeys holding bananas.

responsible for the reversed Stroop effect obtained in experiment 2, then we should obtain significant priming from banana(word) to monkey(word) because of their semantic association. Finally, because we predict a null result—no priming—for the picture task, experiment 4 also included a control naming task in which participants named the *identity* (rather than the color) of the same objects shown in the picture task. The naming task was followed by the same lexical-decision task as that used in the picture and word tasks in order to examine whether picture to word priming was a possible outcome in the picture task. That is, can we establish that there is at least one task in which banana(picture) primes monkey(word)?

5.1 Methods

5.1.1 *Participants.* Experiment 4 included one hundred and twenty-three participants, forty of whom ran in the color-naming of pictures task, forty of whom ran in the color-naming of words task, and twenty-three (two of whom were not included in the analyses because of response key failures) of whom ran in a picture-naming task. All one hundred and three of these participants subsequently ran in the lexical-decision task. Twenty additional participants were run in a preliminary study to obtain conceptually related items for the lexical-decision task.

5.1.2 *Stimuli.* Stimuli for the color-naming tasks were identical to the stimuli used in experiments 1 and 2. Lexical items (36 point Helvetica Bold) for the related condition in the lexical-decision task were obtained by providing the list of color-diagnostic objects (Appendix 1) to participants and for each item having them list two “objects that seem to be strongly associated with the listed objects”. The fifteen responses showing the most consistent associations with a particular object were then selected and used as the related items in the lexical decision task. For the unrelated condition fifteen count nouns were then generated that were not conceptually associated with any items in the color-naming task, but were of approximately equal syllabic length to the fifteen related items. Thirty non-words of approximately equal syllabic length to the thirty related and unrelated words were then generated.

5.1.3 *Design and procedure.* The color-naming phases for both the picture and word groups were identical to the design and procedure used in experiments 1 and 2 respectively. The object-naming phase in the naming group was identical to the picture group except that participants were instructed to verbally identify the object shown on the screen rather than name the color of the object. Following color naming or object naming, all participants, regardless of whether they named colors on pictures, named colors on words, or named objects, ran in a lexical-decision task. On each trial a single letter string appeared and the participant decided as quickly as possible whether the string was a valid English word (all count nouns) or non-word (see figure 3). Of the sixty randomly ordered trials, fifteen were related items, fifteen were unrelated items, and thirty were non-word items. Each item appeared in the middle of the screen and remained displayed until the participant responded by pressing one key for word and a different key for non-word. No accuracy feedback was provided and key-press response times were recorded with an internal Macintosh timer accurate to ± 8 ms.

5.2 Results and discussion

Lexical decisions were 96.5% accurate for participants that ran in the picture color-naming task, 96.7% accurate for those that ran in the word color-naming task, and 97.1% accurate for those that ran in the naming task. For the picture and word groups there were no significant differences in errors across the related, unrelated, and non-word conditions (range 95.2%–97.5%). For the naming group there was a significant difference in errors across the related, unrelated, and non-word conditions (96.5%, 99.4%, and 95.6%, respectively; $F_{2,40} = 5.20$, $p < 0.01$). For all three groups, within-subjects

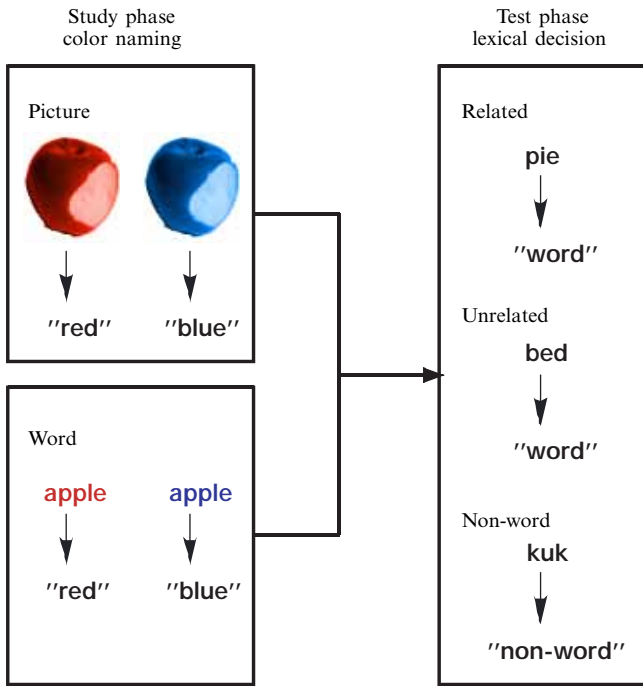


Figure 3. Design of experiment 4. During the study phase, participants performed a color-naming task on either colored pictures or colored words, replicating experiments 1 and 2, respectively. During the test phase, all participants, regardless of whether they had previously seen pictures or words, performed a lexical-decision (word/non-word) task on a series of neutral-colored words. Words presented in the test phase were semantically related to a specific item shown during the study phase, semantically unrelated to any item shown during the study phase, or were non-words. Examples of the stimuli are shown below each condition name in conjunction with the appropriate response for each stimulus item in quotes.

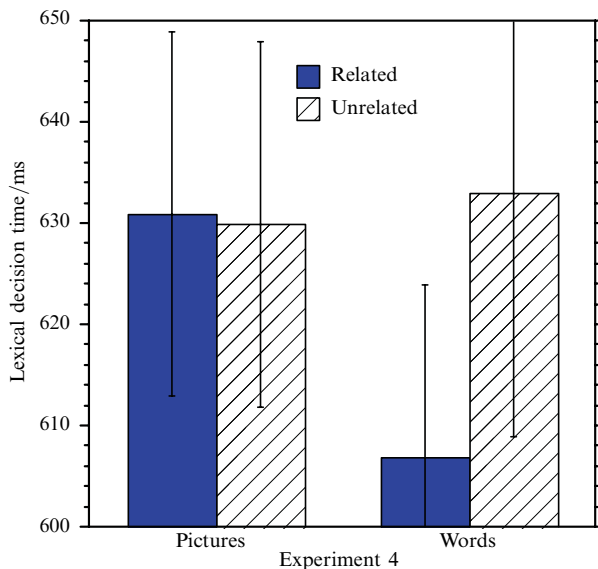


Figure 4. Results of experiment 4 (lexical decision following color naming). Error bars show between-subjects standard error.

one-way ANOVAs on response times with condition as the only factor revealed significant differences between conditions (pictures: $F_{2,78} = 20.49$, $p < 0.001$; words: $F_{2,78} = 28.10$, $p < 0.001$; naming: $F_{2,40} = 24.78$, $p < 0.001$). In all cases, these effects are driven heavily by the much slower response times for non-words as compared to either the related or unrelated conditions. Because the non-word condition is not relevant to the question whether the objects seen during color naming or object naming prime conceptually related items in a lexical decision, the same ANOVA was run including only the related and unrelated conditions. This focused analysis revealed no difference between the related and unrelated conditions (631 ms versus 630 ms, respectively, $F < 1$, for the picture group; a near-significant difference, 607 ms and 633 ms, respectively, $F_{1,39} = 3.62$, $p = 0.06$, for the word group; and a significant difference, 558 ms versus 573 ms, respectively, $F_{1,20} = 4.29$, $p < 0.05$, for the naming group—figure 4). Thus, previously encountering objects specified as words—even when the task is to simply name the color—automatically activates conceptual representations of those objects (hence “banana” facilitated making a lexical decision on “monkey”). In contrast, no hint of lexical priming was observed when participants previously encountered the same objects specified as pictures, suggesting that only visual representations of the objects were recruited in the color-naming task for pictures. That is, color-naming of pictures did not prime related conceptual knowledge about objects, making it unlikely that conceptual representations were activated in this task. Moreover, it is not simply that pictures cannot semantically prime words—when participants named the objects depicted in each picture, priming in the same lexical-decision task was obtained. Thus, we have further evidence that the difference between the typical and atypical color conditions observed in experiment 1 is most plausibly attributed to color–shape associations in *visual* object representations.

6 General discussion

At issue is the nature of color–shape associations in the mental representation of objects. Importantly, such associations may manifest themselves at multiple levels of representation, including the visual, conceptual, and lexical levels. Although it is almost incontrovertible that we *know* that bananas are yellow, it has not been clear whether the *visual memory* for the shape of bananas is associated with the visual memory for the color of bananas. Indeed, this latter question is a point of some controversy in the object-recognition community. On the one hand, Biederman and Ju (1988) concluded that color does not facilitate the recognition of objects when shape alone is sufficient. On the other hand, there is a growing body of evidence that color can play an important role in object recognition, in particular when the color is diagnostic as to the identity of the object (Tanaka and Presnell 1999; Tanaka et al 2001) or the scene (Gegenfurtner and Rieger 2000; Oliva and Schyns 2000). At the same time, several results suggest that color also facilitates object and scene recognition at earlier, sensory levels of visual processing, perhaps by enhancing segmentation (Aginsky and Tarr 2000; Wurm et al 1993; Yip and Sinha 2002).

Relevant to this debate, we interpret our present results as evidence for color being an *intrinsic* component of visual representations of objects (rather than simply aiding segmentation or other precursors to object recognition). This conclusion is based on the Stroop-like effect we obtained in experiment 1 with pictorial stimuli and the failure to obtain conceptual priming with the same pictorial stimuli in experiment 4. This conclusion is further reinforced by experiment 2. Specifically, when orthographic stimuli are used, the effect reverses, to become an ‘anti-Stroop’ effect. Moreover, in experiment 4 we obtain conceptual priming with the same orthographic stimuli. Thus, visual access to associated object shape leads to enhanced color naming, while conceptual and lexical access to associated object shape leads to impaired color naming. Finally, experiment 3

verifies that when methodologies are equated, our materials yield effects consistent with those found in the extant literature on color naming in Stroop tasks.

Of course, these conclusions are open to some re-interpretation. If color–shape associations occur at multiple levels, it is difficult to establish definitively the source of either enhanced or impaired color naming by behavioral measures alone. Specifically, both a picture of an object and the printed name of that object ultimately (and perhaps automatically) lead to conceptual activation (although experiment 4 raises doubts as to the automaticity of this process for pictures). Moreover, object concepts may then in turn activate visual memory or lexical items. We claim, however, that the opposite effects obtained in experiments 1 and 2, and the dissociable effects obtained in experiment 4, point to at least two different loci for color–shape associations—the most plausible sites being the levels accessed *immediately* in each experiment. However, it is also possible that several levels of color–shape association become active simultaneously and that any effect on performance reflects the contribution of multiple levels of access. Critically, in each of our experiments we obtained a highly reliable difference between conditions or reliable priming effects, indicating that, if multiple sources are at work, they appear to contribute in some manner in each case (and differently between experiments or groups).

In summary, there is a growing collection of results supporting a role for color information in object and scene recognition (eg Gegenfurtner and Rieger 2000; Tanaka and Presnell 1999; Yip and Sinha 2002). However, the majority of the studies in this literature have used an object-naming paradigm, which de facto recruits information about objects at multiple levels of representation.

In contrast, we used a color-naming task which does not require access to any information about the shape or the identity of objects. Indeed, observers can respond by simply perceiving the color of the stimulus and then accessing the appropriate color name. We are relying on the automaticity of both shape and word processing to activate the object representation most closely tied to the stimulus modality—pictorial or lexical. We then assume that any impact on the color-naming response must be a consequence of color–shape associations at these immediate levels of access. That we obtain effects in opposite directions across presentation modalities indicates that we have successfully isolated different aspects of the representation of objects. Moreover, that we obtain a reliable effect in all cases indicates that color–shape associations are present by default at multiple levels of object representation. Thus, we can state with some certainty that color is intrinsic to how we learn about, remember, and recognize objects.

Acknowledgments. This research was supported by NSF Award SBR-9615819 to MJT and Grant #1R01EY12691 from the NEI/NIH to MJT and DK. The authors wish to thank Sheila Blumstein, Katherine Bock, William Heindel, and Julie Sedivy for their comments on various aspects of this work, and Anjula Joshi for her help in running the experiments and analyzing the data.

References

- Aginsky V, Tarr M J, 2000 “How are different properties of a scene encoded in visual memory?” *Visual Cognition* **7** 147–162
- Biederman I, Ju G, 1988 “Surface versus edge-based determinants of visual recognition” *Cognitive Psychology* **20** 38–64
- Collins A M, Quillian M R, 1969 “Retrieval time from semantic memory” *Journal of Verbal Learning and Verbal Behavior* **8** 240–247
- Foss D J, Hakes D T, 1978 *Psycholinguistics: An Introduction to the Psychology of Language* (Englewood Cliffs, NJ: Prentice-Hall)
- Gauthier I, Curran T, Curby K M, Collins D, 2003 “Perceptual interference supports a non-modular account of face processing” *Nature Neuroscience* **6** 428–432
- Gegenfurtner K R, Rieger J, 2000 “Sensory and cognitive contributions of color to the perception of natural scenes” *Current Biology* **10** 805–808
- Hayward W G, Williams P, 2000 “Viewpoint dependence and object discriminability” *Psychological Science* **11** 7–12

-
- Klein G S, 1964 "Semantic power measured through the interference of words with color-naming" *American Journal of Psychology* **77** 576–588
- Livingstone M S, Hubel D H, 1984 "Anatomy and physiology of a color system in the primate visual cortex" *Journal of Neuroscience* **4** 309–356
- Ménard-Buteau C, Cavanagh P, 1984 "Localisation de l'interférence forme/couleur au niveau perceptuel dans une tâche de type Stroop avec des stimuli-dessins" *Canadian Journal of Psychology* **38** 421–439
- Meyer D E, Schvaneveldt R W, 1971 "Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations" *Journal of Experimental Psychology* **90** 227–234
- Oliva A, Schyns P G, 2000 "Colored diagnostic blobs mediate scene recognition" *Cognitive Psychology* **41** 176–210
- Price C J, Humphreys G W, 1989 "The effects of surface detail on object categorization and naming" *Quarterly Journal of Experimental Psychology, Human Experimental Psychology* **41** 797–827
- Stroop J R, 1935 "Studies of interference in serial verbal reactions" *Journal of Experimental Psychology* **18** 643–662
- Tanaka J W, Presnell L M, 1999 "Color diagnosticity in object recognition" *Perception & Psychophysics* **61** 1140–1153
- Tanaka J, Weiskopf D, Williams P, 2001 "The role of color in high-level vision" *Trends in Cognitive Sciences* **5** 211–215
- Tanaka K, Saito H, Fukada Y, Moriya M, 1991 "Coding visual images in the inferotemporal cortex of the macaque monkey" *Journal of Neurophysiology* **66** 170–189
- Tarr M J, Bülthoff H H, 1995 "Is human object recognition better described by geon-structural-descriptions or by multiple-views?" *Journal of Experimental Psychology: Human Perception and Performance* **21** 1494–1505
- Williams P, Tanaka J, 2000 "Color effects on the recognition of distinctively shaped objects" unpublished manuscript, Oberlin College
- Wurm L H, Legge G E, Isenberg L M, Luebker A, 1993 "Color improves object recognition in normal and low vision" *Journal of Experimental Psychology: Human Perception and Performance* **19** 899–911
- Yip A W, Sinha P, 2002 "Contribution of color to face recognition" *Perception* **31** 995–1003

Appendix 1

The following list of stimuli was used in all three of the experiments. Each object was rendered both in its typical and an atypical color as listed below.

Object	Typical color	Atypical color
apple	red	blue
apricot	orange	purple
artichoke	green	red
banana	yellow	purple
basketball	orange	green
carrot	orange	green
cheese	yellow	blue
cherries	red	yellow
corn	yellow	red
eggs	white	yellow
garlic	white	purple
grapefruit	yellow	purple
grapes	green	blue
lemon	yellow	red
lime	green	red
lobster	red	yellow
orange	orange	blue
peapod	green	red
pepper	green	blue
pumpkin	orange	red
raspberry	red	purple
rose	red	blue
school bus	yellow	green
sheep	white	green
stove	white	yellow
strawberry	red	green
swan	white	yellow
tennis ball	green	purple
toilet	white	green
traffic cone	orange	yellow

ISSN 0301-0066 (print)

ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 32 2003

www.perceptionweb.com

Conditions of use. This article may be downloaded from the Perception website for personal research by members of subscribing organisations. Authors are entitled to distribute their own article (in printed form or by e-mail) to up to 50 people. This PDF may not be placed on any website (or other online distribution system) without permission of the publisher.