Psychonomic Bulletin & Review
2009, 16 (3), 573-577
doi:10.3758/PBR.16.3.573

Is a bear white in the woods? Parallel representation of implied object color during language comprehension

Louise Connell and Dermot Lynott
University of Manchester, Manchester, England

Color is undeniably important to object representations, but so too is the ability of context to alter the color of an object. The present study examined how implied perceptual information about typical and atypical colors is represented during language comprehension. Participants read sentences that implied a (typical or atypical) color for a target object and then performed a modified Stroop task in which they named the ink color of the target word (typical, atypical, or unrelated). Results showed that color naming was facilitated both when ink color was typical for that object (e.g., *bear* in brown ink) and when it matched the color implied by the previous sentence (e.g., *bear* in white ink following *Joe was excited to see a bear at the North Pole*). These findings suggest that unusual contexts cause people to represent in parallel both typical and scenario-specific perceptual information, and these types of information are discussed in relation to the specialization of perceptual simulations.

Color is an important part of our conceptual representation of objects. Knowledge about color typicality allows us to recognize objects with highly diagnostic colors (e.g., *banana* or *fire engine*) more rapidly than objects with no particular diagnostic color (e.g., *dog* or *lamp*: Tanaka & Presnell, 1999). Indeed, our conceptual knowledge of an object's typical color is more influential in object recognition than is the color actually perceived (Mapelli & Behrmann, 1997; Tanaka & Presnell, 1999). For example, when participants are primed with a picture of a purple apple (i.e., displayed in an atypical color), they are faster to recognize the word *cherry* (which shares the prime's typical color red) than they are the word *blueberry* (which shares the prime's displayed color purple: Joseph & Proffitt, 1996).

However, the presence of context can easily alter conceptual considerations of an object's color. For example, Medin and Shoben (1988) found that people, when asked to compare the color gray with black and with white, considered gray to be more similar to white in the context of *hair*, but more similar to black in the context of *clouds*. Similarly, Halff, Ortony, and Anderson (1976) found that people represented the color red differently for *hair*, *wine*, *flag*, *brick*, and *blood*, considering the color of a red *flag* to be more similar to a red *light* than to a red *wine*. Such context effects are not limited to simple noun–color combinations, but have also been found for larger scenarios. Research in embodied cognition has shown that people represent implied perceptual information during sentence comprehension even though doing so does not facilitate task performance (Connell, 2007; Stanfield & Zwaan, 2001; Zwaan, Stanfield, & Yaxley, 2002). In the case of color, Connell has shown that short-term representations of object color can affect people's ability to recognize objects. For example, when presented with a sentence that implied a particular color for an object (e.g., *Joanne always took milk in her coffee*), followed by a picture (i.e., a cup of coffee), people's speed in verifying that the object had been previously mentioned depended on whether the coffee was shown as milky brown or as straight black.

So what happens if our contextual representation of an object conflicts with our canonical knowledge about its typical state? Theories of embodied (grounded) cognition usually describe color representation as the specialization of a perceptual simulation to include color information (Barsalou, 1999, 2008; Zwaan, 2004). That is, the same neural subsystems that represent color in perception are activated to represent color detail in the conceptualization of an object; specifically, fMRI has shown the same region in the left fusiform gyrus to be implicated in both perceptual and conceptual processing of color (Simmons et al., 2007). However, there has been little discussion of how such specialization might take place if the object simulation is already, by default, specialized with a typical color. For example, we know that tomatoes are usually red, but we may encounter a scenario in which they are green. Which representation—canonical typical or contextual atypical—plays a dominant role? The semantic Stroop task (Klein, 1964; Ménard-Buteau & Cavanagh, 1984; cf. Stroop, 1935) provides an interesting paradigm with which to investigate this question.

L. Connell, louise.connell@manchester.ac.uk

THE PRESENT STUDY

In the present experiment, people were asked to perform a semantic Stroop task that tested whether object-typical and/or context-implied color information was activated during sentence comprehension. Participants were presented with a color-associated object word, such as *tomato* (in typical red, atypical green, or unrelated brown ink), after having just read a context sentence with implicit color information, such as *Jane tasted the tomato when it was ready to eat* or *Jane tasted the tomato before it was ready to eat* (implying either typical or atypical color for the tomato). In Stroop tasks, the effect of context depends on what has been primed: If the target word is primed, it interferes with naming ink color (Warren, 1972), whereas if the color name is primed, it facilitates naming ink color (Cheesman & Merikle, 1986; Kouider & Dupoux, 2004). The design used in the present experiment always primed the target word (because it was always mentioned in the previous sentence), but it primed the name of the ink color according to which color was implied by the context. This manipulation therefore allowed us to examine whether color naming was being facilitated by either object-typical or context-implied color.

The perceptual simulation view holds that language comprehension involves representing implicit perceptual and motor details of a scenario as well as the information explicitly conveyed. For example, in a sentence such as *Joe was excited to see the bear at the North Pole*, the representation of the bear should be specialized with the implied color, white, even though the sentence contains no explicit color information. But will typical brown or contextual white facilitate color naming when the word *bear* is subsequently presented in a semantic Stroop task? Given the importance of color in object representations, it may be the case that typical color is represented regardless of context and is preserved even in the face of specific bias toward an atypical color. If implied color is specialized and represented in parallel with typical color, then responses will be facilitated when the ink color is typical for the given object and when the ink color matches the color just implied (i.e., fastest for *bear* in brown ink when it follows a sentence implying a brown bear, and fastest for *bear* in brown and white ink when it follows a sentence implying a white bear).

Method

Materials. Fifty word–color combinations were created for use in the present experiment. Of these, 30 were test words (object nouns with associated color typicality, forming triads of typical, atypical, and unrelated ink colors—e.g., *bear* in brown, white, and yellow) and 20 were fillers (object nouns with no associated color, each displayed in a single ink color—e.g., *book* in turquoise). All typical and atypical ink colors were sampled from photographs of the relevant objects, meaning that both typical and atypical versions represented possible (natural) colors for that particular object. A pretest of 12 independent raters confirmed that each chosen typical color (e.g., bear–brown) was considered more typical for that object than was its atypical counterpart (e.g., bear–white) at least 75% of the time ($M = 94\%$). Unrelated ink colors represented highly unusual (unnatural) colors for a particular object that would not normally be encountered (e.g., bear–yellow). Color saturation and luminosity were controlled between typical, atypical, unrelated, and filler ink colors.

Forty context sentences were constructed to accompany the target words. Of these, 20 were test sentences (featuring test words; see the Appendix) and 20 were fillers (featuring filler words). Thus, the test sentences formed pairs, with each member of a pair implying either a typical or an atypical color for the same object. Another pretest was conducted to ensure that the test sentences actually implied the intended object color. Pairs of test sentences were separated to form two groups of items, and 24 new participants were randomly assigned to one of the groups. Each sentence was presented along with two line drawings of the target object. One drawing was shaded using the object's typical color, and the other was shaded using its atypical color. Participants were asked to choose, from four forced choice alternatives, whether the sentence was best matched by the first picture, the second picture, both pictures equally, or neither picture. All test sentences that were used implied the intended color (and only the intended color) significantly above the chance level of 25% and equally effectively for typical ($M = 82\%$, $SD = 15\%$) and atypical ($M = 83\%$, $SD = 14\%$) colors.

Participants. Fifty-four native speakers of English with full color vision from the University of Manchester (not involved in pretests) participated for course credit.

Procedure. Participants read instructions describing the experiment that asked them to read each sentence. They were told that a word from the sentence would then appear on-screen and that their task was to name the color of the text—out loud and as quickly as possible—using short color names. Participants were also told that quick responses were important because their response time (RT) was being measured, and that they should read every sentence carefully because their comprehension would be tested at various points during the experiment. A light gray screen background was used throughout the experiment to optimize the visibility of every stimulus color. Each trial began with a left-aligned vertically centered fixation cross that was presented for 1 sec, followed by the presentation of a sentence. When participants pressed the space bar to indicate comprehension, another fixation cross was displayed centrally on-screen for 500 msec, followed by a single word. Participants had to name aloud the ink color of the displayed word as quickly as possible (note that each participant saw each target word in only one possible combination of implied and ink color). In half of all filler trials, a comprehension question (relating to the filler sentence) appeared after color naming and participants indicated their decision by pressing the key, labeled either "yes" (comma key) or "no" (period key). Each participant was required to answer an equal number of "yes" and "no" comprehension questions. A blank screen was displayed for 500 msec as an interstimulus break between trials. The entire procedure took approximately 10 min; this included a practice session to allow participants to become accustomed to the voice-response task.

Design and Analysis. The experiment was a 2 (implied color: typical, atypical) × 3 (ink color: typical, atypical, unrelated) design, with both factors manipulated within participants. We use a convention of implied color–ink color to refer to individual conditions (e.g., a sentence implying that a tomato is red that is followed by the word *tomato* in green ink belongs in the typical–atypical condition).

Data from 3 participants were discarded: 1 for disfluencies on the majority of trials, 1 for responding incorrectly to more than 50% of comprehension questions, and 1 for not being naive to the experimental design, which was revealed during debriefing. Color-naming responses were considered to be correct if the experimenters considered the named color to be a reasonable approximation of the color displayed (e.g., the color for *chameleon* in its atypical ink color was considered correct if it was named as "yellow" or "orange," but not if it was named "red" or "blue"). All responses more than 2.5 standard deviations away from the condition mean were eliminated as outliers (2.7% of data). Since items are not nested under experimental variables (i.e., each target word appeared in all six experimental conditions: typical–typical, typical–atypical, typical–unrelated, atypical–typical, atypical–atypical, atypical–unrelated), an ANOVA by participants is reported (Raaijmakers, Schrijnemakers, & Gremmen, 1999). Effect sizes are reported as generalized eta-squared (η^2_G), which allows direct comparison of within- and between-participants designs (Olejnik & Algina, 2003).

Results and Discussion

Results show that Stroop color naming can be facilitated by both object-typical and context-implied color, which is consistent with the view that the specialization of object color is represented in parallel with typical color. Figure 1 shows the mean correct RTs for the implied color × ink color conditions (typical–typical, $M = 925$, $SD = 242$; typical–atypical, $M = 1{,}028$, $SD = 192$; typical–unrelated, $M = 1{,}055$, $SD = 242$; atypical–typical, $M = 955$, $SD = 224$; atypical–atypical, $M = 967$, $SD = 202$; atypical–unrelated, $M = 1{,}081$, $SD = 272$). Accuracy rates were equivalent across all conditions (93%–97%) and were comparable to filler items (95%).

Color naming was affected by ink color, with naming times overall increasing from typical to atypical to control [$F(2{,}100) = 14.7$, $p < .0001$, $\eta^2_G = .050$], and with an interaction emerging between implied and ink color [$F(2{,}100) = 2.95$, $p = .057$, $\eta^2_G = .008$]. There was no overall effect of implied color [$F(1{,}50) = 0.006$, $p = .94$, $\eta^2_G < .0001$]. A simple-effects analysis showed that the influence of ink color held both when the implied color was typical [$F(2{,}100) = 9.44$, $p < .001$, $\eta^2_G = .059$] and when it was atypical [$F(2{,}100) = 9.48$, $p < .001$, $\eta^2_G = .057$]. When the sentence implied a typical color for an object (e.g., implying that a bear is brown), people were faster to name the ink color when it was object typical (e.g., *bear* in brown) than when it was either atypical (e.g., *bear* in white; planned comparison, $p = .001$) or unrelated (e.g., *bear* in yellow; $p < .001$), with no difference between atypical and unrelated naming times ($p = .419$). On the other hand, when the sentence implied an atypical object color (e.g., implying that a bear is white), both typical and atypical ink colors were named equally quickly ($p = .703$) and faster than were unrelated colors (both ps $= .001$).[1]

It could be argued that color names are simply being primed by terms in the context sentence rather than by the represented colors of objects (e.g., *North Pole* priming "white" and *woods* priming "brown"). To investigate this possibility, we ran a control study in which color-unassociated objects were substituted for the color-associated target objects—for example, *John was excited to see a flag in the woods / at the North Pole*," with the word *flag* subsequently appearing in brown, white, or yellow ink (typical, atypical, or unrelated color, respectively, for the original *bear*). Other substitutions included *piece of litter* for *leaf*, *design* for *steak*, and *flowers* for *tree* (see the Appendix for original sentences). Thirty-six participants (none of whom had taken part in the earlier experiment, and 1 whose data was unusable because of a technical malfunction) were tested using a procedure that was identical to that of the main experiment. A total of 3.5% of responses were excluded as outliers, leaving the following means: typical–typical, $M = 979$, $SD = 164$; typical–atypical, $M = 1{,}034$, $SD = 233$; typical–unrelated, $M = 981$, $SD = 272$; atypical–typical, $M = 992$, $SD = 222$; atypical–atypical, $M = 1{,}023$, $SD = 227$; atypical–unrelated, $M = 1{,}071$, $SD = 357$. All effects disappeared [implied color, $F(1{,}34) = 0.21$, $p = .65$, $\eta^2_G = .001$; ink color, $F(2{,}68) = 0.98$, $p = .38$, $\eta^2_G = .007$], as did the critical interaction [$F(2{,}68) = 0.45$, $p = .64$, $\eta^2_G = .003$]. These results confirm that context sentences in the main experiment act to modulate the color of the target object rather than to activate the target color independently.[2]

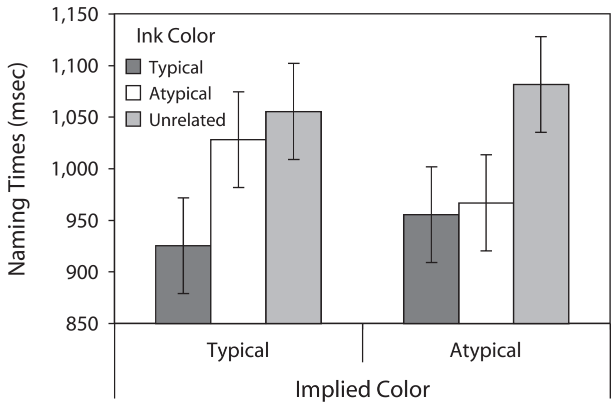

Figure 1. Mean color naming times (in milliseconds) per condition. Error bars represent 95% confidence intervals for multifactorial within-participants designs (Loftus & Masson, 1994).

GENERAL DISCUSSION

In the present study, we investigated how implied information about typical and atypical colors is represented during language comprehension in a novel application of the semantic Stroop paradigm. Results show that color naming was facilitated both when ink color was typical for that object and when it matched the color implied by the previous sentence,[3] suggesting that context-implied information is held in parallel with the more usual, typical information about an object. In other words, when object-typical color is implied, then nothing out of the ordinary has happened and the object retains its usual specialization of the typical color (e.g., *brown bear*), and when object-atypical color is implied, then something unusual is afoot and a parallel specialization of both typical and atypical colors is represented (e.g., *white bear*| *brown bear*).

The effects we reported in the present article are similar to those in other work on situational property priming. For example, McKoon and Ratcliff (1988) showed that a scenario involving a painter searching for a suitable color to depict a ripe tomato activated "red" more than "round," whereas a context about a girl rolling a tomato across the floor showed the reverse pattern. Tabossi (1988) used a time-sensitive cross-modal priming paradigm to determine that the context-appropriate property is activated immediately on processing of the relevant object word (e.g., "yellow" is primed by the scenario *In the light the blond hair of the little girl had the lustre of gold*, but not by *In the shop the artisan shaped with ease the bar of gold*, when tested at the offset of *gold*). However, our study differed from these others in one crucial aspect: We examined different possibilities for a single object property, not different properties for the same object. In other words, a given tomato can easily be both round and red, but the same tomato should not be both red and green. Therefore, context does not just act to reorder or constrain properties' availability according to their salience in the described scenario, but allows mutually exclusive properties to be represented concurrently according to how the object is situated.

There are a couple of implicit assumptions in this account that we will now examine in more detail. First is the issue of time course: How can we be sure that the object color is represented during the processing of the sentence and not at a later point when the color of the target word must be named? For example, a backward integration argument could suggest that people do not represent the color of the bear while reading the sentence; thus, when subsequently naming the "white" ink color, they find it easier to integrate into the North Pole (atypical) context than into the woods (typical) context.[4] Although the backward integration account has been used to question the conclusions of some priming studies regarding noun disambiguation and relative property activation (e.g., Sanford, 1999), it cannot explain the full pattern of the present results. If people only attempt to integrate the given ink color at the point of naming, they should find it easier to integrate "white" into the North Pole context information than they would "brown." This difference does not occur; rather, white and brown were named equally quickly.

The second issue is the nature of the color representation itself: Is color an aspect of the grounded, perceptual simulation of the object as situated in the described scenario, or is it an ungrounded, amodal symbol distinct from perceptual processing areas? In the present study, we presented the effects of context on property typicality within an embodied framework of parallel specialization, although the nature of our research question did not seek to distinguish between the two views. Much research on priming and context effects is agnostic to the nature of the underlying representation, instead focusing on the processes and timing involved in integrating information from language and memory, such as inferencing in discourse (e.g., Graesser, Singer, & Trabasso, 1994; McKoon & Ratcliff, 1992) or lexical ambiguity resolution (e.g., Duffy, Morris, & Rayner, 1988; Giora, 1999). Such behavioral work, including the present study, is concerned with the time course of information access and is compatible with this information being realized in many different forms, including perceptual simulations.

It could be suggested that people are not merely specializing object color in our scenarios, but rather are specializing the subcategorization of the object (e.g., categorizing from generic tomato to specialized ripe tomato or unripe tomato). However, in terms of embodied representations, this is not a matter of particular concern. A simulation of a *green tomato* will also carry other sensorimotor information previously experienced regarding such tomatoes, such as crisp texture and sharp taste. The fact that this tomato may now be labeled an "unripe tomato" is secondary to the specialization itself. In other words, when language comprehenders attempt to specialize an object with context-implied information that is already specialized with prototypical information, they also appear to retain some concurrent representation of the object in its canonical form.

The present findings offer an insight into how the well-documented phenomena of typicality and context effects operate during the comprehension of implied perceptual information. Parallel specialization, in which typical object information is held in mind in the face of contradictory context, offers several advantages to the language comprehender, such as allowing for easy error correction and rapid identification of other (more typical) exemplars. Further research is needed to investigate the implications of such possibilities.

AUTHOR NOTE

Some of this work was carried out when the first author was at Northumbria University. Correspondence concerning this article should be addressed to L. Connell, School of Psychological Sciences, University of Manchester, Oxford Road, Manchester M13 9PL, England (e-mail: louise.connell@manchester.ac.uk).

REFERENCES

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral & Brain Sciences*, **22**, 577-660.

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, **59**, 617-645.

Cheesman, J., & Merikle, P. M. (1986). Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology*, **40**, 343-367.

Connell, L. (2007). Representing object colour in language comprehension. *Cognition*, **102**, 476-485.

Duffy, S. A., Morris, R. K., & Rayner, K. (1988). Lexical ambiguity and fixation times in reading. *Journal of Memory & Language*, **27**, 429-446.

Giora, R. (1999). On the priority of salient meanings: Studies of literal and figurative language. *Journal of Pragmatics*, **31**, 919-929.

Graesser, A. C., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review*, **101**, 371-395.

Halff, H. M., Ortony, A., & Anderson, R. C. (1976). A context-sensitive representation of word meanings. *Memory & Cognition*, **4**, 378-383.

Joseph, J. E., & Proffitt, D. R. (1996). Semantic versus perceptual influences of color in object recognition. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 407-429.

Klein, G. S. (1964). Semantic power measured through the interference of words with color-naming. *American Journal of Psychology*, **77**, 576-588.

Kouider, S., & Dupoux, E. (2004). Partial awareness creates the "illusion" of subliminal semantic priming. *Psychological Science*, **15**, 75-81.

Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review*, **1**, 476-490.

Mapelli, D., & Behrmann, M. (1997). The role of color in object recognition: Evidence from visual agnosia. *Neurocase*, **3**, 237-247.

McKoon, G., & Ratcliff, R. (1988). Contextually relevant aspects of meaning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 331-343.

McKoon, G., & Ratcliff, R. (1992). Inference during reading. *Psychological Review*, **99**, 440-466.

Medin, D., & Shoben, E. (1988). Context and structure in conceptual combinations. *Cognitive Psychology*, **20**, 158-190.

Ménard-Buteau, C., & Cavanagh, P. (1984). Localization of form/color interference at the perceptual level in a Stroop task with drawings. *Canadian Journal of Psychology*, **38**, 421-439.

OLEJNIK, S., & ALGINA, J. (2003). Generalized eta and omega squared statistics: Measures of effect size for some common research designs. *Psychological Methods*, **8**, 434-447.

RAAIJMAKERS, J. G. W., SCHRIJNEMAKERS, J. M. C., & GREMMEN, F. (1999). How to deal with "the language-as-fixed-effect fallacy": Common misconceptions and alternative solutions. *Journal of Memory & Language*, **41**, 416-426.

SANFORD, A. J. (1999). Word meaning and discourse processing: A tutorial review. In S. C. Garrod & M. J. Pickering (Eds.), *Language processing* (pp. 301-334). Hove, U.K.: Psychology Press.

SIMMONS, W. K., RAMJEE, V., BEAUCHAMP, M. S., MCRAE, K., MARTIN, A., & BARSALOU, L. W. (2007). A common neural substrate for perceiving and knowing about color. *Neuropsychologia*, **45**, 2802-2810.

STANFIELD, R. A., & ZWAAN, R. A. (2001). The effect of implied orientation derived from verbal context on picture recognition. *Psychological Science*, **12**, 153-156.

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

TABOSSI, P. (1988). Effects of context on the immediate interpretation of unambiguous nouns. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 153-162.

TANAKA, J. W., & PRESNELL, L. M. (1999). Color diagnosticity in object recognition. *Perception & Psychophysics*, **61**, 1140-1153.

WARREN, R. E. (1972). Stimulus encoding and memory. *Journal of Experimental Psychology*, **94**, 90-100.

ZWAAN, R. A. (2004). The immersed experiencer: Toward an embodied theory of language comprehension. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 44, pp. 35-62). New York: Academic Press.

ZWAAN, R. A., STANFIELD, R. A., & YAXLEY, R. H. (2002). Language comprehenders routinely represent the shapes of objects? *Psychological Science*, **13**, 168-171.

NOTES

1. Since object color was implied rather than directly stated in the context sentences, each participant saw more filler items (containing no color information) than they did test items, and the implied color only matched the ink color in one third of test items; there would be negligible advantage in strategically inferring object colors.

2. We thank an anonymous reviewer for raising this possibility.

3. Unlike in Connell (2007), in which implied color produced interference, the present study found a facilitation effect due to fundamental task differences. In a perceptual task such as object recognition, in which shape has primacy, people can ignore the unhelpful color of a pictured object unless it has been previously activated by a context sentence (Connell, 2007). Naming colors, however, is the *raison d'être* of a linguistic Stroop task, and people are faster to do so when the target color has been previously activated.

4. We thank an anonymous reviewer for raising this possibility.

APPENDIX
Experimental Stimuli

Test sentences with target words and their exact RGB codes (red–green–blue values from 0–255) are presented in the following format: [typical/atypical] implied color → *target word*: typical/atypical/unrelated ink color. Ink color names are approximate.

- The bananas that Mark bought [looked/didn't look] ready to eat → *bananas*: yellow (248–212–14)/green (181–228–36)/red (196–0–51).
- Joe was excited to see a bear [in the woods/at the North Pole] → *bear*: brown (132–75–0)/white (255–255–255)/yellow (248–212–14).
- The teacher pointed to the chameleon lying camouflaged in the [grass/sand] → *chameleon*: green (79–157–37)/yellow (224–159–66)/white (255–255–255).
- Susan liked it when her [granddaughter/grandmother] wore her hair up → *hair*: brown (191–122–55)/gray (170–159–158)/blue (1–204–255).
- Anna found it very [easy/difficult] to spot the lamb in the dark grass → *lamb*: white (255–255–255)/black (81–77–75)/orange (255–103–1).
- Sarah stopped in the woods to pick a leaf off [a tree/the ground] → *leaf*: green (0–130–0)/orange (255–103–1)/gray (170–159–158).
- The children watched the seagulls fly across the sky in the [sunshine/rain] → *sky*: blue (1–204–255)/gray (192–192–192)/green (181–228–36).
- John looked at the steak [on his plate/in the butcher's window] → *steak*: brown (128–63–0)/red (196–0–51)/gray (170–159–158).
- Jane tasted the tomato [when/before] it was ready to eat → *tomato*: red (209–78–13)/green (145–164–54)/brown (132–75–0).
- Paula thought the tree outside her window looked beautiful in the [summer/autumn] → *tree*: green (3–126–0)/orange (228–177–2)/black (132–75–0).

(Manuscript received March 31, 2008;
revision accepted for publication February 12, 2009.)