A JOURNAL OF THE ASSOCIATION FOR PSYCHOLOGICAL SCIENCE aps

Research Article

The Relationship Between Language and the Environment

Information Theory Shows Why We Have Only Three Lightness Terms

Roland Baddeley and David Attewell

University of Bristol

ABSTRACT—*The surface reflectance of objects is highly variable, ranging between 4% for, say, charcoal and 90% for fresh snow. When stimuli are presented simultaneously, people can discriminate hundreds of levels of visual intensity. Despite this, human languages possess a maximum of just three basic terms for describing lightness. In English, these are* white *(or* light*),* black *(or* dark*), and* gray*. Why should this be? Using information theory, combined with estimates of the distribution of reflectances in the natural world and the reliability of lightness recall over time, we show that three lightness terms is the optimal number for describing surface reflectance properties in a modern urban or indoor environment. We also show that only two lightness terms would be required in a forest or rural environment.*

People can discriminate hundreds of levels of visual intensity (Chapanis, 1965), and yet English possesses only three basic terms for describing lightness. English is far from unusual in this paucity of brightness terms. Although the theoretical details of Berlin and Kay's (1969) survey of basic color terms are controversial (Saunders & van Brakel, 1997), it is a robust finding that even the simplest languages (in terms of the number of color terms) have two basic lightness terms, whereas the most complicated have only three. Why should this be? One way of addressing this interesting question is to use information theory (Shannon & Weaver, 1949; Baddeley, 2000).

To apply information theory, we propose that the use of lightness terms should be formalized in terms of the simple language game (Wittgenstein, 1953) shown in Figure 1. This game has two participants (a signaler and a receiver) and takes place in an environment known to both participants (subject to noise due to imperfect memory and perception). The goal of the signaler is simply to view a series of surface reflectances and then, after a delay, describe them to the receiver. Therefore, by describing the viewed surfaces, the signaler reduces the receiver's uncertainty as to their reflectance; information has been transmitted! Information theory allows us to quantify the amount of information required to be transmitted between the signaler and receiver in order to achieve effective communication.

Our analysis of the language game has three stages. In the first stage (Fig. 1a), the signaler samples a surface that he wishes to describe later to the receiver. We assume that reflectance is approximately uncorrelated with behavioral importance, so this reflectance is a random reflectance sample from the given environment (Fig. 1a, schematic graph above the surface). This surface is then illuminated by an unknown illuminant, and the signaler observes the reflected light (Fig. 1a). Because of neuronal noise and, more important, failures of lightness constancy, the signaler's estimate of reflectance based on this observation will have some uncertainty associated with it. This uncertainty is quantified in terms of a probability distribution over reflectances (Fig. 1a, schematic graph on the left), as is usually measured in simultaneous lightness-constancy experiments. So far, the receiver has neither seen nor been informed about the particular surface. However, from prior experience, he does know something about the probability distribution of reflectances in the environment. He does not know the actual distribution, but rather knows a version of it, filtered through his imperfect perception and memory (as illustrated in the right-hand portion of Fig. 1a).

Some time later (Fig. 1b), the signaler meets the receiver and wishes to communicate the surface reflectance of the object he has observed. However, during the intervening period, the sig-

Address correspondence to Roland Baddeley, Department of Experimental Psychology, University of Bristol, 12a, Priory Rd., Bristol, BS8 1TU, United Kingdom, e-mail: roland.baddeley@bristol.ac.uk.

Psychological Science RESEARCH, THEORY, & APPLICATION IN PSYCHOLOGY AND RELATED SCIENCES

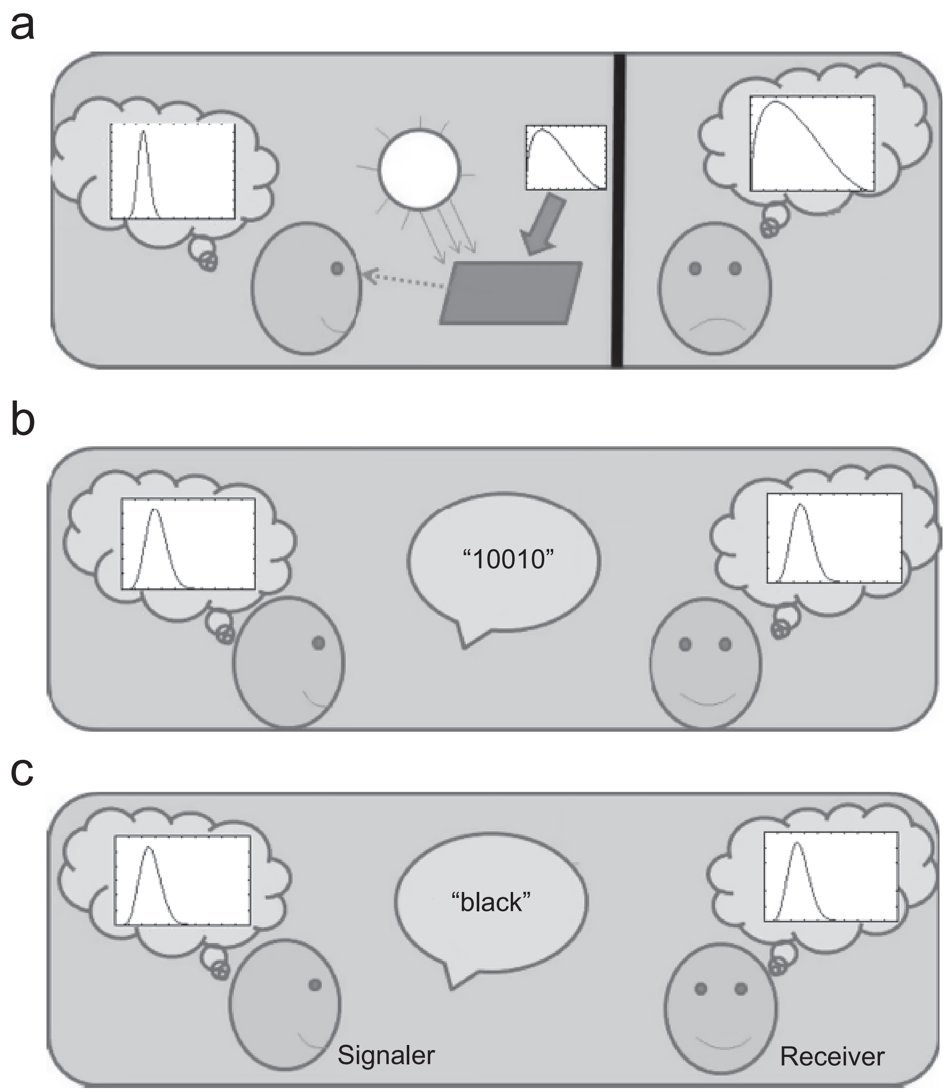

Fig. 1. Analysis of a simple language game to formalize lightness language. The game has two stages. First, the signaler observes a surface (a). Then, the signaler uses binary digits to describe the surface to the receiver (b). We compared the bandwidth the signaler and receiver needed to communicate reflectance via binary digits with the bandwidth they needed to communicate with empirically observed lightness language (c). See the text for details.

naler's uncertainty in the reflectance has increased (left-hand portion of Fig. 1b) due to both low-level (contrast and luminance adaptation) and high-level (decay in memory accuracy) factors. It is this uncertain knowledge of reflectance that he has available to communicate to the receiver. This is done by transmitting information through a communications channel (e.g., "10010" in Fig. 1b). Information theory quantifies the information in this channel in terms of the equivalent number of binary yes/no questions (bits) that would be required in order to reduce the receiver's uncertainty (i.e., right-hand portion of Fig. 1a) to the same level as that of the signaler (right-hand portion of Fig. 1b).

This would be fine if people communicated using binary digits, but in order to determine how the number of bits transmitted relates to the number of lightness symbols we use (e.g., "black" in Fig. 1c), we need to know the capacity of these symbols to communicate. This depends not only on the number of symbols (lightness terms), but also on their relative probability of use. Communication is maximized when all symbols are used with equal probability, and minimized when one symbol is used almost exclusively. Analysis of this last stage therefore involves quantifying the probability distribution of lightness terms in real-world languages, and hence their ability to communicate information.

So far, what we have is a general framework for understanding communication about physical characteristics of the world. It is made specific to lightness communication by the characterization of the three key distributions: the distribution of reflectances in the world, the distribution of uncertainty in reflectance introduced over time by memory and perception, and the distribution of word-use frequency in brightness language.

The most important of these is the distribution of reflectance signals in the world (Fig. 1a). This enters our game as the distribution of surfaces from which the signaler samples, and as the basis of the receiver's model of the world before he receives information from the signaler. The effect of this distribution is simple: The more varied (higher entropy) the environment, the more terms needed to describe it.

The second distribution we need to characterize is the uncertainty in the signaler's estimate of reflectance when he comes to communicate the surface lightness. This uncertainty will vary as a function of the reflectance level, and again this has a simple effect: The less certain the signaler is about the reflectance after a delay, the less information he has to communicate to the receiver and the fewer the number of terms needed to communicate lightness.

The last distribution is that of lightness term use in everyday language. Again, the pattern is simple: More uniform use of terms will result in greater capacity to signal effectively and, therefore, fewer terms will be required to achieve a given level of information transmission.

We describe the estimation of the distribution of surface reflectances in different environments, a simple experiment to measure the uncertainty in people's recall of reflectances as a function of reflectance level, and the information theory techniques required to calculate the number of lightness terms that should be stable within a given environment. The results that we obtained were related to what is known about the relationship between language and the world: an area where color (and its brightness subset) has often been used to explore various theoretical proposals. We found that lightness language is well matched to the world.

METHOD

Here, we describe details of the procedures and calculations involved in estimating the brightness information that needs to be conveyed about the environment, the distribution of reflectances in the environment, the uncertainty in our internal representation of a real-world reflectance due to memory and failures of lightness constancy, and the entropy of lightness terms in English.

Information theory is a method for quantifying the amount of information (in bits) one quantity provides about another. In this case, we quantified how much information, $I(R'|R)$, one language user's memory of the surface reflectance, R', gives another language user, regarding an object's actual reflectance, R. This places an upper limit on the channel capacity (and hence how many symbols would be required) if the observer wished to communicate about this reflectance: If each symbol is used equally often, the number of symbols required is $2^{I(R'|R)}$ (we also considered the effects of using symbols with nonequal frequencies).

$I(R'|R)$ is based on the difference between two measures that represent the uncertainty in the internal estimate of reflectance before and after a reflectance has been viewed (i.e., the uncertainty shown in the right-hand portions of Figs. 1a and 1b):

$$I(R'|R) = H(R') - H(R'|R). \quad (1)$$

The first term on the right is the entropy associated with the internal representation of reflectance, that is, the uncertainty in R' prior to viewing a reflectance in the environment, R:

$$H(R') = -\int_0^1 P(R')\log(P(R'))dR',$$

where $P(R)$ is the distribution of reflectances in the environment, and $P(R')$ is the distribution of the internal representation of reflectances:

$$P(R') = \int_0^1 P(R'|R)P(R)dR.$$

The second term in Equation 1 is the conditional entropy of the internal representation of reflectance, that is, the average uncertainty in R' given that R has been viewed:

$$H(R'|R) = -\int_0^1 P(R)\int_0^1 P(R'|R)\log(P(R'|R))dRdR'.$$

To evaluate Equation 1, we therefore needed estimates of both $P(R)$, the distribution of reflectances for a given environment, and $P(R'|R)$, the distribution of the internal representation of reflectance, conditional on a reflectance having been viewed. Previously, we used biological survey techniques to estimate the distribution of reflectances in a range of environments, including modern indoor and woodland environments (Attewell & Baddeley, 2007). We found, after comparing a range of low-dimensional parameterizations of these distributions, that they were best fitted by beta distributions; for example, $P(R)$ = beta(α, β), with $\alpha = 1.29$, $\beta = 2.3$, for a modern indoor environment; $\alpha = 1.35$, $\beta = 10.72$, for an urban exterior environment; and $\alpha = 1.91$, $\beta = 22.6$, for the woodland environment. In other words, the distribution of reflectances in an indoor environment is far from uniform, being on average far darker (mean = 33%) and far less variable (SD = 0.22). Reflectances in a woodland environment are darker still (mean reflectance = 8%) and less varied (SD = 0.06), with reflectances in an outdoor urban environment being between these extremes (mean reflectance = 11%, SD = 0.11).

We next explored the uncertainty in the reflectance of a surface after it has been viewed by the receiver: $P(R'|R)$. This uncertainty is due to noise within perception (imperfect lightness constancy) and imperfect memory. We estimated $P(R'|R)$ by first getting 36 subjects to memorize eight different target objects (crosses,

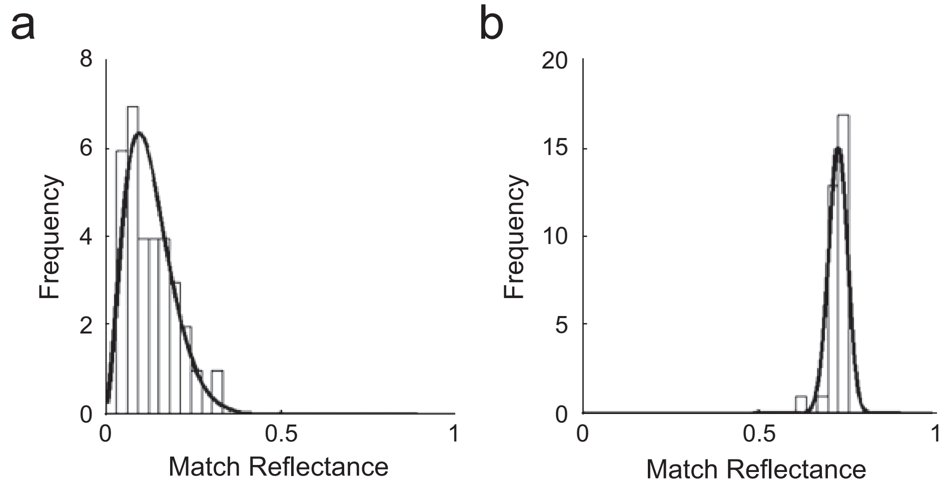

Fig. 2. The distribution of subjects' reflectance responses for two test levels of reflectance: (a) 17% and (b) 71%. The histogram of responses and the best-fitting (maximum likelihood) beta distributions are shown. The beta distributions provide a good characterization of the data.

squares, circles, etc.). Target items were presented in two blocks of four items. Four items is within the limits of short-term memory capacity (Cowan, 2001). Each object had a different shape, name, and surface reflectance (chosen at intervals of 0.09 between 0.08 and 0.71). After 7 min, a collection of 24 objects of the same shape, but with reflectances ranging from 0.05 to 0.74 in steps of 0.03, was presented. Subjects were asked to point to the object that had the most similar reflectance to the target. Averaging over subjects allowed us to estimate the conditional distribution of the group's internal representation. For each of the levels of reflectance tested, the conditional distribution of the group's internal representation was again well characterized by a beta distribution (see Fig. 2 for the distribution of responses for two different levels of target reflectances). Because we were examining communication between individuals, we were interested in this average distribution over subjects. By allowing subjects to point rather than verbally describe the reflectance, we removed the communication constraints imposed by language, because pointing has a large channel capacity; in this case, $\log_2(24) = 4.5$ bits. The variance in recalled luminance values across individuals has been shown to change relatively little when the retention interval is increased from 15 s to 15 min, and to not change at all when the interval is increased from 15 min to 24 hr (Perez-Carpinell, Baldovi, de Fez, & Castro, 1998). Therefore, although we only use a single retention interval of 7 min, our method should produce estimates of the distribution of recalled reflectances that are robust over the interval durations between viewing and subsequently describing a surface encountered in the real world.

Our method allowed us to estimate $P(R'|R)$ for the eight reflectance levels tested. However, to calculate the channel capacity, we needed to have a noise estimate for all possible reflectances. To provide a beta distribution estimate for $P(R'|R)$ of all values of R, we used cubic spline interpolation to interpolate over both $\alpha/(\alpha + \beta)$ and $\alpha + \beta$. This spline interpolation is shown in Figure 2; as can be seen, the spline provides a good characterization of the data.

These methods allowed us to calculate the channel capacity of lightness language, and therefore how many brightness terms would be required if all terms were used with equal frequency. However, this pattern of term frequency is rarely seen in natural languages. To obtain estimates of the information transmitted by natural languages, with which we could compare our estimates of channel capacity, we calculated the entropy of three synthetic lightness languages based on English. The first was a two-term language with only the terms *light* and *white* (which we treat as synonyms) and *dark* and *black*. The second was a three-term language that included the term *grey/gray*. Finally, we used a five-term language that also contained the separate composite terms *light grey/gray* and *dark grey/gray*. The entropy of a language tells us how much information, on average, is given by each term. This entropy is calculated from the probability distribution of term use using the following equation:

$$H(T) = \sum p(t)\log_2 \frac{1}{p(t)},$$

where $H(T)$ is the information transmitted by the term used to describe a lightness, and $p(t)$ is the probability of using term t.

The probability of use of the various simple and compound brightness terms within these languages was estimated using word frequency data from four different sources. First, we used the Internet as a corpus, and used the number of search returns for the individual terms provided by three Internet search engines (MSN, Google, and Ask) as estimates of term frequency. No constraints other than the terms themselves were placed on the searches. To avoid counting instances where *grey/gray* occurred as part of a compound term when determining the number of instances of the lone basic color term *grey/gray*, the sum of the number of returns for *light grey/gray* and *dark grey/gray* were subtracted from the number of returns for *grey/gray* alone. To check the validity of this (to our knowledge) novel approach, we also obtained word frequencies from a more traditional analysis of the British National Corpus of written and spoken English (Leech, Rayson, & Wilson, 2001).

Ideally, we would have used word frequencies taken from natural two-term languages, but because such languages are generally unwritten, there is no literary corpus from which to derive the frequencies.

TABLE 1

The Information and Number of Lightness Terms Required to Describe Four Environments

Variable	Environment: Artificial uniform	Modern interior	Modern urban exterior	Deciduous woodland
Information transmitted (bits)	2.7	1.48	0.96	0.72
No. of lightness terms	6.5	2.78	1.91	1.65

Note. The number of lightness terms represents the minimum required if all lightness terms are used with equal frequency.

RESULTS

Figure 2 shows the distribution of subjects' remembered surface reflectances for two of the eight reflectance standards used. As can be seen, for both of these situations (and all others tested), the distribution of responses (the subjects' uncertainty in the reflectance) is well approximated by a beta distribution.

Our experimental results give a parametric characterization of the uncertainty for the eight levels we tested, but our calculation requires an estimate for all reflectance levels. To estimate the reflectance levels, we interpolated the parameters of our beta functions as shown in Figure 3. As can be seen, especially over the important range between 0.1 and 0.5, where most reflectances occur, the interpolation is good.

This distribution of uncertainty, together with the distribution of reflectances within the environment, allowed us to calculate the amount of reflectance information available to the signaler when describing a surface, and hence the number of lightness terms required to transmit this information, assuming equal frequency of term use (see Method). These results are given in Table 1.

Our last task was to estimate the entropy of lightness terms in English for simulated languages containing two lightness terms (*black* and *white*), three lightness terms (*black*, *white*, and *grey/gray*), and five lightness terms (*black*, *white*, *grey/gray*, *light grey/gray*, and *dark grey/gray*). To do this, we used the frequencies of hits for these terms obtained from three different Internet search engines (Google, MSN, and Ask), and frequency of use within the British National Corpus. The estimated entropies are given in Table 2. The word-frequency data in Table 3 show clearly that the term probabilities obtained using Internet searches are in close agreement with those obtained using a traditional corpus analysis.

Table 2 shows that adding the qualifiers *light* and *dark* to form compound lightness terms has little effect on the entropy of a five-term language over that of a three-term language. We use the composite terms so infrequently (see Table 3) that they

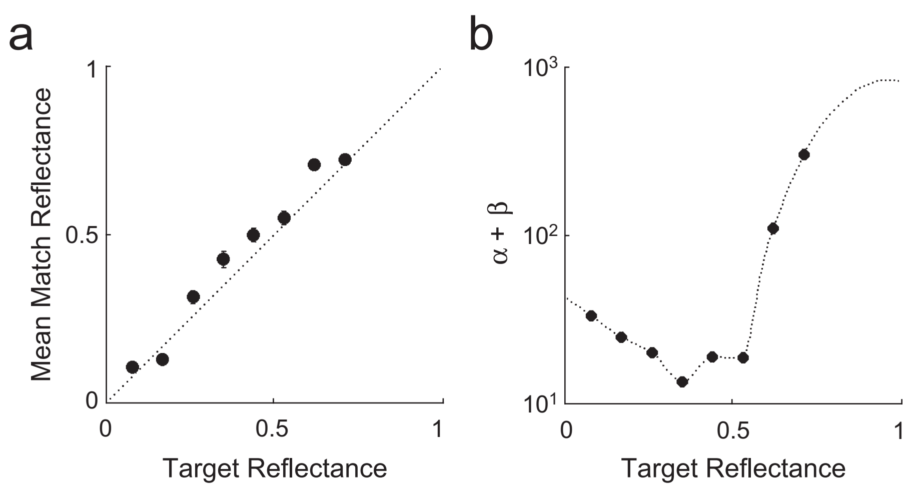

Fig. 3. Mean (a) and spread (b) of the beta distribution as a function of target reflectance. The mean of the beta distribution was calculated as $\alpha/(\alpha + \beta)$. The spread of the beta distribution was calculated as $\alpha + \beta$. A cubic spline was used to estimate these values for intermediate reflectances. Although the high reflectance end was sparsely sampled, this made little difference to our calculations because environmental reflectances are predominantly low (below 50%).

TABLE 2

Entropy of Lightness Terms in Two-, Three-, and Five-Term Languages

	Entropy (bits)			
No. of terms	Google	MSN	Ask	British National Corpus
Two terms	0.972	0.889	1.000	0.999
Three terms	1.357	1.143	1.411	1.367
Five terms	1.383	1.157	1.440	1.474

Note. Lightness-term data were obtained from three Internet search engines (Google, MSN, and Ask) and the British National Corpus of written and spoken English (Leech, Rayson, & Wilson, 2001). Entropies were calculated using the word-frequency data given in Table 3.

contribute little to the average information transmitted by a lightness term within a five-term language.

DISCUSSION

We used estimates of the distribution of reflectances from four environments (uniform, modern interior, urban exterior, and woodland) to characterize the reflectance signal. We used a simple delayed match-to-sample experiment, in which subjects pointed to the best estimate of the reflectance after a short delay, to estimate the level of uncertainty due to inaccurate memory and perception. Given these estimates, it is straightforward to calculate the amount of lightness information an observer, within a given environment, can convey about reflectance. This gives us an estimate of the number of lightness terms that may be used consistently and therefore remain stable with a language. We then compared this to the estimated channel capacity of synthetic languages, based on English, which had two, three, or five terms. The match between the observed and predicted channel capacity was very close (Fig. 4), predicting both the number of terms and how this number changes depending on the environment. This finding suggests that the number of lightness terms within a language represents an optimal solution to the problem of describing the variation in reflectances encountered within the visual environment while ensuring that these descriptions can be made consistently and accurately.

Fig. 4. The lightness information present in natural languages as a function of environment. For comparison, the entropy (channel capacity) for lightness ''languages'' with two, three, and five lightness terms (see Method) is shown.

Why do we get this result of two to three terms, depending on environment? For an artificial world in which all reflectances between 0% and 100% are equally likely (i.e., the reflectance distribution with the maximum variability), the channel capacity would be 2.7 bits (6.5 or effectively 7 lightness terms). This result is very similar to a number of previous studies that have investigated information transmission in artificial environments. For brightness, Eriksen (1954) found that information about reflectances drawn from a uniform distribution requires 2.3 bits (5 terms). This earlier phase of information theory research, which ignored the distribution of signals in the real world, is very elegantly summarized in Miller's (1956) famous "magic number seven" article.

TABLE 3

Relative Probability of Lightness Terms in Two-, Three-, and Five-Term Languages

	Relative term probabilities											
	Google			MSN			Ask			British National Corpus		
Lightness term	Two terms	Three terms	Five terms	Two terms	Three terms	Five terms	Two terms	Three terms	Five terms	Two terms	Three terms	Five terms
Black	.402	.360	.358	.693	.657	.655	.492	.432	.431	.522	.470	.462
White	.598	.534	.532	.307	.290	.290	.508	.447	.445	.478	.430	.423
Gray	—	.106	.106	—	.053	.053	—	.121	.120	—	.100	.098
Dark gray	—	—	.002	—	—	.001	—	—	.002	—	—	.098
Light gray	—	—	.002	—	—	.001	—	—	.002	—	—	.002

Note. Term-frequency data from which the probabilities are derived were obtained from four different sources: the number of search returns from three Internet search engines (Google, MSN, and Ask) and the British National Corpus of written and spoken English (Leech, Rayson, & Wilson, 2001).

The real world is far less variable than a uniform distribution (and far darker). When moving from the laboratory to the real world, our analysis shows that information transmission goes down from requiring 6.5 lightness terms to 2 or 3 terms (from 2.7 bits to 0.72–1.48 bits, depending on the environment). This finding suggests that part of the reason that we have few lightness terms is that the world is not very rich in terms of reflectances. Even in the most variable environment measured in this study, and perhaps one of the most variable environments regularly encountered by humans in terms of lightness, the modern interior, only three terms are required. However, we might expect artists or paint-shop workers, who work with and communicate about a wider, more uniform range of reflectances, to have up to 7 stable lightness terms within their work-specific vocabularies.

Although composite lightness terms such as *light gray* and *dark gray* do exist within natural languages, we found that they are used so infrequently as to have little effect on the average lightness information transmitted by a term (Table 3). This result is in agreement with Biggam's finding in her study "Grey in Old English" that "compound color terms are a rarity in greyness" (1998, p. 313), and suggest that, when describing lightness within an everyday real-world environment, three terms represents a real upper limit to the number of terms required to transmit the lightness information available.

Of the three real-world environments whose lightness distributions we use in this study, only the modern indoor environment contains sufficient lightness information to require three lightness terms. However, the term *gray* is likely to have been in place alongside *black* and *white* in Old English when it arrived in England with the Anglo-Saxons in the 5th century AD (Biggam, 1998). Although we do not suggest that Anglo-Saxons inhabited anything akin to a modern interior environment, it is clear that they did not live within a purely woodland environment, which, as we have seen, can be described using only two lightness terms. Painted walls, dyed cloth, and polished metalwork were common (Westwood, 2008). Indeed, the literary corpus from which our understanding of their lightness language is derived was created by scholars and scribes who, in the very act of recording, were exposed to white vellum or parchment, thereby greatly increasing the variation in lightness experienced day to day. It is our belief that this variation would have been sufficient to necessitate the use of a third lightness term (*gray*) in order to transmit the lightness information available.

A second reason why we have a small number of terms is our high level of uncertainty in the reflectance of a previously viewed surface, even after a short period of time. We used a recall interval of 7 min, but previous research (for instance, Perez-Carpinell et al., 1998) showed that recall of a target's luminance becomes drastically worse after a period as short as 15 s. Although performance in simultaneous luminance matching can be very good, even after a very small temporal interval, luminance-recall performance decays rapidly. This, we believe, may be due to failures of lightness constancy. When viewing two surfaces at the same time, it is reasonable to assume that both the illumination and the statistics of the surrounds are similar. When two targets are widely separated in time, it is not reasonable to assume similar illumination, adaptation level, or background statistics. All these contribute to achieving lightness constancy, so the accuracy or inaccuracy of memory may be matched to the constancy of the luminance signal (i.e., its accuracy in representing reflectance).

Lastly, the requirements for the method presented here are very basic and easy to obtain. All that is required is a characteristic of the world that fulfills four criteria. First, it must be behaviorally relevant. Second, it must be described verbally. Third, we must be able to estimate its probability distribution in the world. Finally, we must be able to characterize the uncertainty in the remembered value of this quantity. Examples could include the verbal description of angles or of people's ages. In the latter case, the distribution of the signals would be very different for different people: In particular, young mothers would experience distribution of "signals" (ages of subjects) that was much more biased to young ages than, say, elderly men. The number and type of linguistic labels for these two groups would therefore be predicted to be different. This seems, at least intuitively, very plausible. One area that would be of particular interest is to apply the method to general color words. Here, the difficulty is not theoretical but practical. The number of data points needed to estimate a probability distribution scales with the power of its dimensionality, and color is three dimensional. This means that the number of examples from the natural world required to obtain a good estimate of the distribution would be much larger. Also, characterizing the noise for all colors would require a lot of work; however, with approximations, this should be possible. It would be intriguing to see whether we find the same number of categories as identified originally by Berlin and Kay (1969).

In conclusion, we suggest that natural languages have a maximum of only three terms for lightness because it represents the optimal number for describing our uninteresting visual environment (in terms of reflectances), given the uncertainties introduced by memory and perception.

REFERENCES

Attewell, D., & Baddeley, R.J. (2007). The distribution of reflectances within the visual environment. *Vision Research*, *47*, 548–554.

Baddeley, R.J. (2000). Introductory information theory and the brain. In R.J. Baddeley, P.J. Hancock, & P. Foldian (Eds.), *Information theory and the brain* (pp. 1–20). Cambridge, England: Cambridge University Press.

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

Biggam, C.P. (1998). *Grey in Old English: An interdisciplinary semantic study*. London: Runetree Press.

Chapanis, A. (1965). Color names for color space. *American Scientist*, *53*, 327–346.

Cowan, N. (2001). The magical number four in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 87–185.

Eriksen, C.W. (1954). *Multidimensional stimulus differences and accuracy of discrimination* (U.S. Air Force, Wright Air Development Center Technical Report No. 54-165). Washington, DC: U.S. Air Force.

Leech, G., Rayson, P., & Wilson, A. (2001). *Word frequencies in written and spoken English: Based on the British National Corpus*. London: Longman.

Miller, G.A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*, 81–97.

Perez-Carpinell, J., Baldovi, R., de Fez, M.D., & Castro, J. (1998). Color memory matching: Time effect and other factors. *Color Research & Application*, *23*, 234–247.

Saunders, B.A.C., & van Brakel, J. (1997). Are there nontrivial constraints on colour categorization? *Behavioral and Brain Sciences*, *20*, 167–179.

Shannon, C.E., & Weaver, W. (1949). *The mathematical theory of communication*. Chicago: University of Illinois Press.

Westwood, J.O. (2008). *Celtic and Anglo-Saxon art and ornament*. New York: Dover.

Wittgenstein, L. (1953). *Philosophical investigations*. Oxford, England: Blackwell.

(Received 3/14/08; Revision accepted 1/11/09)