tion. A repertoire of language is apparently much larger than that of motion. Therefore, it may provide a new source for the sensory-motor coordination. We thus expect that the diversity of a color category is enhanced by using language, if language can generate more novel sensory-motor coordination of colour experience.

Finally, it is more interesting to study the category for intermediate colors (tones), not the colors themselves, because tones are more subtle than colors and are mixtures of prototypical colors. Without having to worry about such intermediate colors, communication may become easier. Interestingly, however, by using communication the intermediate colors become diverse. In other words, we feel that not only the convergence of color categories, but the increasing diversity of intermediate colors are both caused by social and linguistic communication. This is what we consider to be a creative role of communication beyond sensory-motor coordination.

Sharing perceptually grounded categories in uniform and nonuniform populations

Kimberly A. Jameson
Institute for Mathematical Behavioral Sciences, University of California – Irvine, Irvine, CA 92697-5100 and The Center for Research in Language, University of California–San Diego, La Jolla, CA 92093-0526.
kjameson@uci.edu
http://aris.ss.uci.edu/cogsci/personnel/kjameson/kjameson.html

Abstract: Steels & Belpaeme's (S&B) procedure does not model much of the important variation that occurs across human color categorizers. Human perceptual variation and its corollary consequences impact real-world color categorization. Because of this, investigators with the primary aim of understanding color categorization and naming across cultures should exercise some caution extending these findings to explain how different human societies lexicalize color appearance space.

Steels & Belpaeme (S&B) clearly state that their simulated "perceptually grounded color categories" do not strive to model human categorical representation, but, rather, are practical models of color categorization by artificial embodied agents, or "robots." Their aim is to clarify the conditions under which robot category repertoires make feasible robot communication with human color categorizers. Their approach to color categorization is powerful and refreshingly comprehensive. They synthesize different constraints and contributing factors – biology, psychology, and culture – that are typically pitted against each other in the color cognition and categorization literature.

Despite the authors' statements concerning the work's limited applicability to the behavior of living categorizers, readers of the article are likely to extend these findings to other forms of color categorization phenomena, including human color categorization across cultures. For this reason, discussion is needed of the findings' implications for category processing within and between humans.

S&B state "the agent's architecture is intended to model what we know today about human colour perception, categorization, and naming." (sect. 2.1). Within a population "all agents are assumed to have exactly the same perceptual process." (sect. 2.3.1). And agents base their categorization task decisions (sect. 2.4) on sensory input "So" from a standard uniform perceptual representation (i.e., CIE [Commission Internationale de l'Éclairage, or International Commission of Illumination] L*a*b*). S&B vary spectral distributions that are sampled, but the stimuli are always first converted from spectral distributions to CIE tristimulus values, and then to CIE L*a*b* values before agents engage in any categorization games. Thus agents make all categorization decisions on CIE translations of spectra, rather than on actual spectra (eq. 5). S&B make the following related assumptions:

(A) All agents embody a CIE standard model.

(B) All agents in a population uniformly replicate the same perceptual process.

It should be noted that the above-mentioned details are important considerations if one is seeking an independent neural network or computational modeling verification of human color categories (Cf., Yendrikhovskij 2001a), because these details impose consequences on the artificial network that will influence the network solutions obtained.

Assumption (A) places undeniable constraints on the shared category solutions obtainable by a population of agents. It fixes the dimensionality, metameric class relations and the gamut of the stimulus space to be equivalent with the CIE standard observer. This is a good idea when engineering a robot that strictly aims to perceive spectra with a standardized human eye. Under such circumstances one would minimally expect agents to categorize stimuli in ways compatible with humans, because much of stimulus structure (i.e., metamer equivalences and relations, dimensional structure, and the perceptual gamut of the space) is preprogrammed.[1] Indeed, predefining metameric equivalences alone is enough to establish how agents lump spectra into equivalence classes. This preprogrammed lumping of spectra, however, would not match how most other terrestrial species sort spectral stimuli. Thus, the CIE network-input settings used by S&B differ from others that could be programmed into the agents – say, a standard observer model for male spider monkeys, turtles, or the honeybee – all of which would predictably produce category repertoires fundamentally different from those the authors obtained.[2] Consequently, in addition to S&B's clear disclaimers about the generalizability of the processes by which their robots establish category repertoires, those interested in human color categorization should note that it is also wrong to infer that their networked agents replicate human color categorization behaviors because it is purportedly an optimal species-independent way of categorizing the available terrestrial spectra (this view about optimality appears in the cognitive literature, e.g., Shepard 1994; 1997).

Because assumption (A) predetermines the dimensionality, metamers, and gamut of the categorized color space, a generalization of (A) would also accommodate subspaces of natural categories from agents possessing *fewer* dimensions and *restricted* gamuts. However, such subspaces bring S&B the additional challenge of modeling different metametric class equivalences.[3] Such network modeling adjustments for (A) would be an important step towards modeling human categorization, and bear on assumption (B).

Assumption (B) limits extending S&B's findings to human color categorization, because real human groups that develop and share categorical repertoires are not comprised of individuals with uniform perceptual processing, or uniform color processing expertise. Indeed, relatively minor perceptual variation could significantly impact the network solutions that S&B report. First, considering just perceptual processing variation across agents (i.e., differences of dichromacy and anomalous trichromacy compared to trichromacy), such subgroup processing could impact convergence rates and robustness of solutions under the simplest situations (i.e., learning without language, sect. 3). Discrimination game outcomes could vary for dichromat agents, compared with anomalous-trichromat or trichromat agents, in accord with the observation that "even small variations in colour perception . . . drive . . . colour categories to diverging results" (sect. 5.1). Interactions between actual dichromats and trichromats suggest that perceptual variation effects could extend beyond single agent processing to learning with language scenarios (sect. 4) and guessing game outcomes, making plausible the idea that agent perceptual variation could effect robustness and variance of a population's category repertoire, and, in turn, indices of discrimination success and number of converged on categories.

Human dichromats occur at different rates across ethnolinguistic societies, and, with varying degrees of effectiveness, communicate using trichromat-based lexical categories for which they have no perceptual distinctions (e.g., Jameson & Hurvich 1978;

Shepard & Cooper 1992). In one society where rod monochromacy commonly occurs in the population, color normal individuals share a pragmatic categorical repertoire with achromatopes who perceive a "colorless" world (Sacks 1997). In other societies, other complexities arise during processes wherein perceivers learn through social interaction to use normative linguistic codes despite perceptual differences that could undermine the code's meaning (Jameson 2005a; 2005b; Jameson et al. 2001). Thus, within populations, variation in perceptually correlated knowledge is integral to the cognitive side of learning and sharing a color repertoire, but such human variation runs counter to Assumption (B).

Addressing both (A) and (B) as suggested here would permit S&B to make useful comparisons between perceptually grounded categories shared by uniform populations and those shared by nonuniform populations.

NOTES

1. This seems to work against the suggestion that "artificial agents might end up with a quite different categorical repertoire compared to . . . human beings" (sect. 1).

2. Just as S&B demonstrate different sets of "chromatic distributions . . . do not lead to categories that are similar . . ." (sect. 5.1), so too would very different category solutions arise if initially agents were given a honey-bee observer model, and these category solutions would almost certainly bear little resemblance to the category solutions they found using their agent populations.

3. Just as dichromats are accommodated by the CIE standard observer model, but have different known metameric class relations.

Seeing and talking: Whorf wouldn't be satisfied

Boris Kotchoubey

Institute of Medical Psychology and Behavioral Neurobiology, University of Tübingen, 72074 Tübingen, Germany.
boris.kotchoubey@uni-tuebingen.de
http://www.uni-tuebingen.de/medizinischepsychologie/stuff/

Abstract: Although Steeles & Belpaeme's (S&B) results may be useful for development of technical devices, their significance for behavioral sciences is very limited. This is because the question the authors asked was "Why do people use similar words in a similar way?" rather than "How can similar words stand for similar experience?" The main problem is not shared word usage, but shared references.

> Polonius: What do you read, my lord?
> Hamlet: Words, words, words.
> —*Hamlet*, Act II, Scene II

The clarity with which the target article is written makes the critique easier. The main goal is formulated from the very beginning: To explore how colour words "may become sufficiently shared among the members of a population" (sect. 1) so that if I say "red" everybody can select a red (and not a yellow) object from a presented set. Moreover, Steels & Belpaeme (S&B) make no secret that this "goal is entirely practical . . . to design . . . robots that are able to do this task." (sect. 1) Though I am not an expert in robotics, it appears that the authors attained substantial progress in approaching their goal.

The question is, however, whether this pragmatic approach can shed light on the real mechanisms in question. I agree that the study can contribute to "designing agents that are able to develop a repertoire of . . . categories that is sufficiently shared to allow communication" (sect. 6). But I doubt that "these results are relevant to . . . an audience of cognitive scientists" (sect. 6) who are interested in the psychology of colour perception. Although the authors admit that "the artificial agents might end up with a quite different categorical repertoire compared to . . . human beings," (sect. 1) they miss a much worse peril, that their agents come to categories very similar to human categories (thereby creating the illusion of relevance), but using processing means that have nothing in common with those used by human brains.

S&B suggest that their data support the Sapir–Whorf thesis on the dependence of colour perception on language. This thesis has been formulated in rather ambitious terms, for instance, by Sapir: "We see and hear and otherwise experience very largely as we do because the language habits of our community predispose certain choices of interpretation" (cited by Whorf 1962, p. 134), or by Whorf's commentator S. Chase: "Speakers of different languages see the Cosmos differently" (ibid, p. x). Particularly, Whorf emphasised the importance, not only of verbal categories, but rather of the syntax of different languages (e.g., tenses, subject–predicate structure, use of plurals and singulars, etc.), in organisation of our basic mechanisms of perceiving and conceiving of the world.

This expected relationship to the very structure of colour experience is lacking in the target article. Not sharing perception (e.g., the fact that you see red where I also see it) but sharing word usage is the problem the entire study is pivoted around. By the way, colour may not be the best case for study interaction between sensory and cognitive factors because the sensory information can only be obtained with central vision (there are no cones on the periphery) and high luminance (cones do not work in twilight), hence one may state that we see most objects grey most of the time. But the main point is that mere agreement in verbal behavior does not prove the agents' similarity in their "segmentation of the face of nature" (Whorf 1962, p. 241).

Of course, we cannot really know another person's sensory qualia (e.g., the qualium of redness), but we can approach this knowledge by using a broad range of methods, beyond categorisation and naming. And probably the most reliable result obtained to date is that if we vary tasks, conditions, instructions, cue availability, and so forth, so also varies the role of language as a determinant of behavior. Thus, the long-assumed effect of language spatial terms, such as "on the left of" or "to the north of,"on space perception proved to be the effect of available spatial cues. Natural peoples, when tested in their natural conditions, use significantly more objective (allocentric) spatial cues than Europeans (Dutch or English) tested in the lab. Also English-speaking people, without changing their mother language, use more allocentric cues when tested outdoors as compared to being in a closed room with blinds pulled down (Li & Gleitman 2002). The availability of potentially useful information appears, therefore, to exert a stronger effect on space perception than the language itself.

Turning back to colours, the data are not very different. For example, most European languages have one basic term for blue, whereas Russian has two; a popular Russian children's song listing "the seven colors of the rainbow" mentions light-blue and dark-blue as two completely different colours, the latter being close, but not identical, to purple. Nevertheless, being presented with a large number of green and blue colour tones, Russian and English subjects did not differ in their classification; particularly, Russians did not tend to group dark and light blue separately (Davies & Corbett 1997). There is no evidence that English speakers are unable to distinguish those hues that Russian speakers do.

Kay and Kempton (1984) developed colour triads, such as one containing two green colours and one blue. One of the green colours (Green 1) was separated from the other green (Green 2) by a larger number of just noticeable differences than from Blue. When asked to choose the stimulus that looked least like the other two, subjects chose Blue. However, when asked to compare stimuli pairwise, they found Green 1 and Green 2 more different than Green 1 and Blue. The issue may be even more complicated because neuropsychological data indicate that a patient who performed like controls in this experiment (and who, therefore, could distinguish between classification and similarity judgment) was nonetheless unable to classify colours according to their names. His sorting was based on superficial perceptual similarity (Robertson et al. 1999). This may indicate that not only the presence of