

of the phenotype. Finding such a correlation is neither an end unto itself, nor is it really an answer; rather, it is, at best, only a first methodological step that can be used to manipulate and explore the developmental process at hand (Nijhout 1990).

The bottom line here is that the S&B's nativist model, despite its success, is not a good model of genetic development of colour categorization. The proof for this lies in the very evolutionary dynamics that the authors' model provides. Focal perception is a complex process. It is not only determined by genetics, but by development, neurophysiological constraints, as well as experience. The highly rapid and diverse evolutionary dynamics observed in S&B's experiment is hardly, if ever, possible for complex neurophysiological functions (Bowmaker 1998; Lickliter & Honeycutt 2003; Schlichting & Pigliucci 1998; Surridge et al. 2003; Worden 1995). This in itself casts doubts on the appropriateness of the model.

In sum, I applaud S&B's efforts to model category sharing. By the same token, I believe that they need to come up with a better model for the genetics of colour categorization. I encourage them to do so.

ACKNOWLEDGMENT

I thank Joe Hilferty for discussions of this article. Portions of this article have been extracted from a submitted manuscript by Oscar Vilarroya and Joseph Hilferty "The gene out of the bottle" and from the communications presented by the same authors at the 8th International Cognitive Linguistics Conference (Logroño, Spain, July 2003).

Learning colour words is slow: A cross-situational learning account

Paul Vogt^{a,b} and Andrew D. M. Smith^a

^aLanguage Evolution and Computation Research Unit, School of Philosophy, Psychology, and Language Sciences, University of Edinburgh, Edinburgh EH8 9LL, United Kingdom; ^bInduction of Linguistic Knowledge, Computational Linguistics, Tilburg University, 5000 LE Tilburg, The Netherlands. paulv@ling.ed.ac.uk
<http://www.ling.ed.ac.uk/~paulv> andrew@ling.ed.ac.uk
<http://www.ling.ed.ac.uk/~andrew>

Abstract Research into child language reveals that it takes a long time for children to learn the correct mapping of colour words. Steels & Belpaeme's (S&B's) guessing game, however, models fast learning of words. We discuss computational studies based on cross-situational learning, which yield results that are more consistent with the empirical child language data than those obtained by S&B.

Steels & Belpaeme (S&B) have successfully shown how computational modeling can contribute greatly to the study of the evolution of language and cognition. S&B have – in our opinion correctly – decided to write their article from an engineer's point of view. We feel, however, that their model of linguistic communication would have been more realistic, and therefore the results they obtained more robust, if they had used a model of acquiring colour categories through multiple contexts.

S&B model the communication between agents using the *guessing game* model, which is, in itself, not unreasonable. Their claim, however, that this game is "equivalent" to colour chip naming experiments carried out by anthropologists (sect. 2.4.2), is not justified, in our opinion. The guessing game is primarily a model of learning through corrective feedback, whereas colour chip naming experiments consist of an anthropologist (A) asking an informant (B) to point out, on a chip set, the focal colour of a colour term from B's language. There are three important differences between the anthropological experiments and the guessing game. First, B is not doing any learning, in fact, A is learning about B's representation of colour and about B's language. Second, A does not correct B's responses or provide any feedback about them. Finally, there is no negotiation between A and B about what the words should refer to.

This positive feedback loop between the choice of which words to use and their success in communication is the main learning mechanism in the guessing game. Indeed, S&B claim that the feedback loop is a *necessary* requirement for cultural language development (sect. 5, condition 1), although in fact it is widely accepted that children receive little, if any, corrective feedback while learning words (Bloom 2000, but see Chouinard & Clark 2003, for an alternative account). In computational simulations of lexicon creation and learning, similar to those presented by S&B, we have shown that agents using a *cross-situational statistical learner* (a variant of Siskind's, 1996, cross-situational learner) can successfully develop a shared vocabulary of grounded word meanings *without* corrective feedback (Smith 2003; Vogt 2004). In our model, as in guessing games, hearers have to infer what speakers are referring to, but unlike in guessing games, the agents do *not* have any way of verifying the effectiveness of their attempts at communication. Instead, the agents use covariances to learn a mapping between words and categories based on the cooccurrence of words and potential referents across multiple situations.

Although young children do learn to relate colour terms to colours, it takes them a considerable length of time to find the appropriate mappings (e.g., Andrick & Tager-Flusberg 1986; Sandhofer & Smith 2001). For example, it has been estimated that, on average, children required over 1,000 trials to learn the three basic colour terms "red," "green," and "yellow" (Rice 1980, cited in Sandhofer & Smith 2001). Sandhofer and Smith suggest that children go through different stages in learning colour words: First they appear to learn that colour terms relate to the domain of colour, and only then can they actually learn the correct mapping. This has also been observed by Andrick and Tager-Flusberg (1986), who additionally suggest that children find it difficult to learn the boundaries of colour categories, thus slowing down the learning of colour words. Research into child lexical acquisition is, of course, dominated by the problem of referential indeterminacy, and many constraints have been suggested to explain how children reduce indeterminacy (see, e.g., Bloom 2000). Very few of these accounts, however, allow for the fact that children hear words in multiple different contexts, and can use this to determine the intended reference. Recent empirical research, indeed, shows that a cross-situational model of learning provides a robust account of lexical acquisition in general, and of the acquisition of adjectives, including colour categories, in particular. Houston-Price et al. (2003) suggest that the children in their study used cross-situational learning to disambiguate word reference, even though their experiments were designed with attentional cues. In addition, Mather and Schafer (2004) show that children can learn the reference of nouns by exploiting covariations across multiple contexts. Akhtar and Montague (1999) demonstrate that children use cross-situational learning to discover the meanings of novel adjectives. Klibanoff and Waxman (2000), furthermore, provide empirical support for their proposal that adjectival categories are learned cross-situationally, within the context of basic level categories.

A comparison of the guessing game and a cross-situational statistical learner, using computational simulations, has shown that, in the guessing game, coherence in production between agents is considerably higher and that learning is much faster (Vogt & Coumans 2003). This means that agents using cross-situational statistical learning have considerable difficulties in arriving at a shared lexicon, although in the end they manage to overcome them. Note, however, that cross-situational statistical learning improves when: agents' semantic categories are similar (Smith 2003); learners assume mutual exclusivity (Smith 2005); and the context size is relatively small (Smith & Vogt 2004). This slower rate of acquisition is thus consistent with the empirical evidence that children learn colour words relatively slowly. Importantly, as yet unpublished studies have shown that the category variance among agents in the cross-situational learner tends to be much higher than that seen from the guessing games. This suggests that negotiating category boundaries in the cross-situational learner is more

difficult, which could confirm Andrick and Tager-Flusberg's (1986) finding.

S&B have presented a model of learning colour words that is fast and based on corrective feedback. Research on child lexicon acquisition suggests, however, that colour categories are actually acquired slowly and through cross-situational learning. If cross-situational learning is, indeed, a more plausible model than the guessing game, then the results achieved by S&B may no longer hold for their account of cultural learning.

Interindividual variation in human color categories: Evidence against strong influence of language

Thomas Wachtler

Department of Physics/Neurophysics, Philipps-University, 35039 Marburg, Germany. thomas.wachtler@physik.uni-marburg.de
<http://neuro.physik.uni-marburg.de/~wachtler>

Abstract: With respect to human color categories, Steels & Belpaeme's (S&B's) simulations over-emphasize the possible influence of language. In humans, color processing is the result of a long evolutionary process in which categories developed without language. Common principles of color processing lead to similar color categories, but interindividual variation in color categories exists. Even color-deficiencies, causing large differences in color categories, remain inconspicuous in everyday life, thereby contradicting the hypothesis that language could play a role in color category formation.

The main focus of Steels & Belpaeme's (S&B) study is category formation in artificial agents and the role language could play in this process. Beyond that, they consider the possible relevance of language for color categories in humans. Neither issue seems to be adequately addressed by the simulations.

First, under the conditions assumed by S&B, it is almost trivial that language would make categories of the simulated artificial agents more similar. S&B specify their stimuli in a homogeneous color space. Because there are no constraints or dissipative mechanisms, *any* kind of coupling will increase the similarity between categories, and eventually lead to identical categories. S&B introduce such coupling in their simulations with color genes, language, or nonuniform stimulus distributions. Not surprisingly, in all of these cases, categories of different agents become similar. Even for the simple case of artificial agents, however, success of the "sharing by language" strategy requires that communication corresponding to "guessing games" would occur with fairly high frequency.

Second, for the case of human color categories, the scenarios considered by S&B are similarly inappropriate. They ignore, for example, the properties and constraints of neural processing and representation in the visual system. There are strong nonlinear mechanisms, such as the division in On- and Off-pathways, which effectively segregates color space into categorical half-spaces. Chromatic preferences of color-selective neurons tend to cluster, both at precortical stages (e.g., Derrington et al. 1984) and in the visual cortex (De Valois et al. 2000; Kiper et al. 1997; Komatsu et al. 1992; Lennie et al. 1990; Wachtler et al. 2003). In other words, not all chromaticities are equal. So far, the exact relation between coding at early stages of the visual system and perceptual categories is still unclear (see e.g., the comments on Saunders & van Brakel 1997; for a recent discussion see Valberg 2001). Nevertheless, nonuniform distribution of color preferences places constraints on category formation. Similarities between the properties of neurons in the visual system and efficient codes for natural colors (Caywood et al. 2004; Lee et al. 2002; Wachtler et al. 2001) further indicate that color vision is adapted to the statistics of natural chromatic signals, which implies shared categories.

The corresponding genetic coupling of color categories is not

realized by "color genes," but rather by the genes that control the development and function of the visual system. These genes evolved over many millions of years, and evolutionary success was not determined by successful communication, but by efficient processing of visual information, probably including such important tasks as image segmentation and the finding of food (e.g., Mollon 1989). Experimental evidence for shared color categories has been found in other species, as well, such as chimpanzees (Matsuno et al. 2004) or even flies (Troje 1993).

Despite the common processing principles underlying human color vision, there are considerable interindividual differences in the prereceptoral, receptor, and postreceptor stages of visual processing. As a result, for example, the loci of unique hues are broadly distributed (e.g., Webster et al. 2000). How does this variability compare to the results of S&B's simulations? S&B fail to specify how similar categories have to be in order to be "sufficiently shared." In any case, with respect to human color vision, S&B's ideal of "complete" sharing is not realistic.

Variation in color vision is most striking in "color-blind" subjects. In dichromats, such as protanopes or deuteranopes, one type of cone photoreceptor is entirely missing. Interestingly, despite their receptor color space of reduced dimensionality, dichromats use the same basic color terms as trichromats when asked to describe their color percepts (Boynton & Scheibner 1967). However, their category regions in color space differ considerably from those of trichromats (Wachtler 2004). This is not surprising, because certain colors belonging to different categories of trichromats, such as trichromats' reds and greens, are indistinguishable for dichromats. Nevertheless, dichromats seem to possess perceptual categories corresponding to those of trichromats, and they seem to achieve them by dividing their reduced color space using both spectral composition and luminance (Boynton & Scheibner 1967; Jameson & Hurvich 1978; Wachtler et al. 2004).

Several lines of evidence indicate that the color categories of dichromats revealed by color naming reflect perceptual categories. For example, dichromats claim that "red," "green," "blue," and "yellow" constitute unique and different percepts. Furthermore, dichromats consistently report a "red" contribution both in short-wavelength and in long-wavelength stimuli, asserting that the "red" is of the same perceptual quality in both cases (Wachtler et al. 2004).

The color naming behavior of color-deficient observers suggests that language plays a role in the acquisition of the lexicon of color names, but does not influence perceptual categories. Just like color-normals, dichromats have to learn the words to name their percepts. Given that the structure of their color space is different, they cannot achieve a perfect match, so they assign those names to their perceptual categories that constitute the best possible match to those of trichromats.

It is impossible for dichromats to have the same categories as color-normal trichromats. If communication about color would be as crucial as S&B suggest, dichromats would be lost in continuous frustration. No matter how long they would learn, it would be impossible for them to adjust their categories accordingly. Luckily, however, color naming plays only a marginal role in everyday life, situations that require accurate communications of color are extremely rare. Many color-deficient individuals are not even aware about their condition until their first color-vision test. Thus, language, although important in establishing a consistent lexicon for our color categories, is too weak a link to influence perceptual categories.

ACKNOWLEDGMENT

I thank Rainer Hertel for inspiring discussions and critical reading of the manuscript.