

PUBLISHED VERSION

Christensen, Wayne David; Tommasi, Luca.

Color categories in biological evolution: broadening the palette, *Behavioral and Brain Sciences*, 2005; 28 (4):492-493.

Copyright © 2005 Cambridge University Press

PERMISSIONS

<http://journals.cambridge.org/action/stream?pageId=4088&level=2#4408>

The right to post the definitive version of the contribution as published at Cambridge Journals Online (in PDF or HTML form) in the Institutional Repository of the institution in which they worked at the time the paper was first submitted, or (for appropriate journals) in PubMed Central or UK PubMed Central, no sooner than one year after first publication of the paper in the journal, subject to file availability and provided the posting includes a prominent statement of the full bibliographical details, a copyright notice in the name of the copyright holder (Cambridge University Press or the sponsoring Society, as appropriate), and a link to the online edition of the journal at Cambridge Journals Online. Inclusion of this definitive version after one year in Institutional Repositories outside of the institution in which the contributor worked at the time the paper was first submitted will be subject to the additional permission of Cambridge University Press (not to be unreasonably withheld).

10th December 2010

<http://hdl.handle.net/2440/38105>

How to learn a conceptual space

Antonio Chella

Dipartimento di Ingegneria Informatica, Università di Palermo, 90128
Palermo, Italy. chella@unipa.it <http://www.csai.unipa.it/chella/>

Abstract: The experiments proposed in the article by Steels & Belpaeme (S&B) can be considered as a starting point toward a general methodology for the automatic learning of conceptual spaces.

In recent years, several frameworks for cognitive robotics have been proposed that take into account a level that is intermediate between the “subsymbolic” low level, directly linked to the external sensors, and the “linguistic” high level, oriented toward symbolic inferences.

A cognitive intermediate level of this kind has been proposed by Gärdenfors (2000). Different from other proposals, Gärdenfors introduces an intermediate level, based on “conceptual spaces,” with a precise geometric structure. Briefly, a conceptual space is a metric space whose dimensions are related to the quantities processed by the agent sensors. Examples of dimensions could be color, pitch, volume, and spatial coordinates. Dimensions do not depend on any specific linguistic description: A generic conceptual space comes before any symbolic-propositional characterization of cognitive phenomena.

A point in a conceptual space is the epistemologically primitive perceptive element at the considered level of analysis. Chella et al. (1997; 2000) describe a robot vision system based on conceptual spaces in which each point corresponds to a *geon*-like 3-dimensional geometric primitive (Biederman 1985) perceived by the robot. Therefore, the perceived objects, like the agent itself, other agents, the surrounding obstacles, and so on, are all reconstructed by means of *geons*, and they all correspond to suitable sets of points in the agent’s conceptual space. A related conceptual space has been proposed by Edelman (1999), which also proposes an implementation based on Radial Basis Functions (RBF) neural networks. Song and Bruza (2003) adopted a conceptual space framework for information retrieval applications, and Aisbett and Gibbon (2001a) propose a suitable conceptual space for clinical diagnosis applications. From a theoretical point of view, Gärdenfors and Williams (2001) discuss the conceptual space approach for generating nonmonotonic logic inferences, and Chella et al. (2004) discuss conceptual spaces in the framework of the anchoring problem in robotics. Balkenius (1998) proposes a more realistic implementation of a conceptual space, from an empiric point of view, by a set of RBF units, and Aisbett and Gibbon (2001b) discuss a related implementation based on voltage maps.

One of the problems with all of the previously cited approaches is that the structure of the adopted conceptual spaces are a priori defined by the designer according to the addressed problem, in the sense that the designer has to define how many axes are necessary for a correct representation of the problem at hand, what is the meaning of the axes and the corresponding type and range of values, what are the separable and the integral dimensions, and so on. No general methodology has been adopted or proposed to allow the machine to inductively learn a conceptual space, with the exception of the multidimensional scaling algorithm (Shepard 1962a; 1962b) proposed by Gärdenfors, which is generally not suitable for real world robotic applications.

Analyzing the article by Steels and Belpaeme (S&B) from the point of view of the conceptual space theory, the described agents effectively build a conceptual space in order to represent the perceived colors. A “category,” implemented by a RBF neural network, identifies a subspace of integral dimensions of colors, because each RBF unit defines a color subdimension, whereas different categories correspond to separable subspaces of colors. Therefore, the color conceptual space of the agent is generated by the union of all the subspaces of integral dimensions of colors corresponding to all the agent categories. The agent inner representation of a color is therefore given by the collection of the re-

sponses of all the RBF units built by the agent, that is, by the components of the conceptual space dimensions, in agreement with the conceptual space theory. It should be noted that each color subspace is implemented by a RBF neural network, along the lines of the approaches by Edelman and by Balkenius.

The new and important point brought forth in the S&B experiments is that the agent conceptual space is not defined a priori by the system designer, but it is learned by the agent itself according to its inner and external constraints, as fully described in the target article. Therefore, the strategy adopted by S&B is effectively able to address the previously described problem of how to learn a conceptual space. Interestingly, the conceptual space is generated not only by means of the agent perceptions, but also by the linguistic interactions among agents, that is, by means of the agent actions.

Along this line, it would be interesting to investigate the possibility for an agent to have more powerful representation capabilities that allow the agent to infer the conceptual spaces of other agents, through, for example, a sort of higher-order guessing game. In this way, the problem of sharing categories among populations could be correctly addressed, in the sense that an agent belonging to a population A_x may build an inner representation of the conceptual space of another agent belonging to a population A_y to acquire all the needed capabilities to “translate” its own color categories to the color categories of the other agent.

In conclusion, the S&B article is a seminal starting point for the investigation of a general methodology for inferential learning of conceptual spaces from an agent’s external perceptions, its inner and external constraints, and its actions.

Color categories in biological evolution: Broadening the palette

Wayne D. Christensen^a and Luca Tommasi^b

^aDepartment of Philosophy, University Kwazulu-Natal, Durban 4041, South Africa; ^bKonrad Lorenz Institute for Evolution & Cognition Research, Altenberg 3422, Austria. Christensen@ukzn.ac.za
<http://www.kli.ac.at/personal/christensen/homepage.html>
luca.tommasi@kli.ac.at
<http://www.kli.ac.at/institute-b.html?personal/tommasi>

Abstract: The general structure of Steels & Belpaeme’s (S&B’s) central premise is appealing. Theoretical stances that focus on one type of mechanism miss the fact that multiple mechanisms acting in concert can provide convergent constraints for a more robust capacity than any individual mechanism might achieve acting in isolation. However, highlighting the significance of complex constraint interactions raises the possibility that some of the relevant constraints may have been left out of S&B’s own models. Although abstract modeling can help clarify issues, it also runs the risk of oversimplification and misframing. A more subtle implication of the significance of interacting constraints is that it calls for a close relationship between theoretical and empirical research.

Steels & Belpaeme’s (S&B’s) study attempts to combine research objectives for robotics and human science. But, although using human communication as a model may be a useful starting point for robotics, the radical differences in physical constraints between robots and humans makes it unclear how much overlap there need be between the two areas. The evolution of human communication abilities occurred in a specific biological context, with perceptual, motor, cognitive, social, and ecological constraints that don’t apply to robots. Exotic abilities like direct sharing of perceptual information are possible for robots, and ultimately the most effective robot communication systems may be no more similar to human verbal communication than human communication is to that of honeybees or dolphins. This is not to suggest that there will be no important commonalities, but rather to point out that divergent specific constraints can generate very different possibilities.

The study is at a sufficiently high level of abstraction that such differences enter the picture only minimally. In taking inspiration from categorization and naming by humans, S&B only focus on the fact that humans are capable of open ended generation of socially shared names, without attempting to model any specific biological or psychological mechanisms that may be involved. Such a strategy offers the potential for generality, but it also faces challenges in demonstrating that the results will be robust against departures from S&B's assumptions, and that the models are informative despite the lack of realism. In their conclusion, S&B suggest that additional realism would only obscure the dynamics, but this very much depends on assuming that unmodeled constraints do not contribute to the dynamics.

Their primary hypothesis is that "embodiment and statistical regularity of the environment is not enough to achieve sufficient sharing for communication and that cultural constraints also play a role." However, this sounds suspiciously like its addressing an ill-formed problem: What counts as "sufficient sharing" and "not enough" may well be sensitive to a variety of factors and vary in different contexts. For example, focal color categorization is present in birds and might have been selected for because it simplifies the cognitive demands of discriminating multiple items in an array, be they landmarks in the environment (Tommasi & Vallortigara 2004), potential mates (Bennett et al. 1997), or the incentive value of different types of food (Gamberale-Stille & Tullberg 2001). Thus, task complexity is one dimension in which departures from S&B's assumptions could have a significant impact: The complex task demands faced by birds may have favored focal color categorization independent of any social referencing constraints. The general problem is that it is very hard to know in advance where and how such interrelations arise, and hence, it can be hard to evaluate whether an abstract model has aptly represented the issues.

Comparative research provides a strategy for disentangling some of these kinds of complexities. In the case of color categorization, birds and mammals are an informative contrast, given the evolutionary radiation that separated these amniote groups from a common ancestor, and the highly sophisticated visual abilities of birds (Güntürkün 2000; Vallortigara 2004). Color perception is more refined in avian species from the level of retinal photoreceptors, because the presence of double cones, oil droplets, and tetrachromacy provide for earlier color-opponency processes than those found in mammals (Vorobyev et al. 1998).

Birds are thus endowed with the structural and functional features necessary to perceive, discriminate, and generalize color stimuli. Selective pressures undoubtedly shaped color spaces in the direction of those sensory aspects that are ecologically relevant to the species. Some behavioral responses (e.g., feeding behavior) are genetically biased in the direction of specific colors (Roper & Marples 1997), but, not surprisingly, the development of chromatic perception is dependent on the statistical structure of the colors experienced in the environment, because rearing newly hatched birds in abnormally colored environments results in alterations of the spectral range to which the birds respond when compared with control animals in color discrimination tests (Miklósi et al. 2002).

As noted, birds have been shown not only to discriminate and generalize colors, but also to categorize the color continuum in discrete regions centered around focal points, as found in humans (Jones et al. 2001). Even more strikingly, birds have exhibited spontaneous emergence of vocalizations akin to color naming. Manabe et al. (1995) trained budgerigars to emit a high pitch call in case of the presentation of one color and to emit a low pitch call in case of the presentation of another color. Once this association was learned, spontaneous differential vocalizations were observed in response to forms when some new association of forms to colors was being established, as if the birds were anticipating the presentation of a color by its learned vocal label. Research on parrot's chattering has provided evidence of color referencing mediated by vocal communication and apparently equally depending on

both parrot-parrot and parrot-human observational learning (Pepperberg & Wilcox 2000). Birds share a basic neural architecture that is substantially different from that of mammals, and yet their evolution independently achieved functions strikingly similar to humans, with the potential for categorization of color in referential communication.

Several points can be drawn from this. Comparing similar traits across diverse phyla is a useful strategy for casting light on evolutionary processes and biological mechanisms, and can help disentangle the general from the specific. However, the differences also make such comparisons fraught; analogies (e.g., describing a particular bird behavior as "naming") must be drawn very carefully. Similar problems of interpretation face theoretical modeling research. If a bird species ever evolves language that involves color naming, do S&B's results imply that the color categories must be socially shaped? We suggest that this is far from clear because it isn't clear how well the assumptions of S&B's models will map to the constraints operative in the particular evolutionary process. Given the difficulty of predicting a priori which constraints, or even which kinds of constraints, may prove relevant in evolutionary and cultural processes, there is reason to try to develop a close coupling between empirical and theoretical research so that the respective strengths and weaknesses can be balanced against each other. In arguing this, we are not seeking to dismiss S&B's work. It is an elegantly structured study that may provide a robust modeling platform for much productive theoretical exploration. But closer empirical links will help its development.

In the beginning: Word or deed?

Stephen J. Cowley

Department of Psychology, University of Hertfordshire, Hatfield, Hertfordshire AL10 9AB, United Kingdom and School of Psychology, University of KwaZulu-Natal, South Africa. s.j.cowley@herts.ac.uk

Abstract: Emphasizing that agents gain from culture-based patterns, I consider the etiology of meaning. Since the simulations show that "shared categories" are not based in learning, I challenge Steels & Belpaeme's (S&B's) folk view of language. Instead, I stress that meaning uses indexicals to set off a replicator process. Finally, I suggest that memetic patterns – not words – are the grounding of language.

Using remarkable simulations, Steels & Belpaeme (S&B) show why autonomous robots can gain from sensitivity to culture-based patterns. These can be used to supplement categories grounded in embodiment and, as a result, actions can be better coordinated. The simulations thus illuminate "how the memetic evolution of language and meaning is possible" (sect. 4.3). In this commentary, stressing that agents use indexical signs, I focus on the etiology of meaning. Language itself, I suggest, may depend on how grounded categories interact with memetic patterns or indexical signs.

Although "sharing" develops in genetic simulations, as for color, this may be gross. Equally, as with herring gull chicks, it may depend on "relational signs" that arise in the niche (Tinbergen 1953). Further, the simulations show that shared categories will not arise from individual learning. Mapping a word-form to a color is, for this reason, beyond autonomous devices that lack sensitivity to culture-based patterns. It becomes possible, however, provided that an encultured pattern is consistently coupled with what sensors can detect. Given learning, coupling can give a population grounded relational categories.

S&B draw on a folk view of language. Taking shared categories for granted, they assume that a lexicon maps words onto meanings. Accordingly, they adopt what has been called the "fundamental assumption of linguistics" – the view that, "in every speech community, some utterances are alike in form and meaning" (Bloomfield 1935, p. 78; for critique, see Love 2003). In spite of