must automatically switch between incremental unsupervised clustering, when predictive feedback is unavailable, and supervised clustering, when predictive feedback is available, without a loss of memory stability, to learn categories that are sensitive to both environmental statistics and cultural constraints. ARTMAP algorithms achieved this goal by combining top-down matching and attention with vigilance or sensitivity control that together maximize generalization while minimizing predictive error using a patented concept called match tracking (Carpenter & Grossberg 1991). In this regard, S&B discuss distinguishing between edible and nonedible mushrooms, notably the need to distinguish "fine shades of orange" (sect. 2.2) in mushroom databases. However, they do not simulate classical mushroom databases (Schlimmer 1987). ARTMAP has been benchmarked on the mushroom database with a 99.8% accuracy during on-line learning (Carpenter et al. 1991).

S&B also note the importance of studying "categorization and naming by humans" (sect. 1), but do not do model human performance. ARTMAP has simulated the set of thirty human categorization experiments, called the 5–4 category structure (Smith & Minda 2000), which is a standard benchmark for human categorization (Ersoy et al. 2002). Whereas traditional cognitive models can fit these data, they do so without learning the categories and without describing underlying brain dynamics. ARTMAP learns the categories and fits the data at least as well as cognitive models, and also proposes how to settle the classical exemplar/prototype debate concerning whether exemplars or prototypes are stored in memory. ARTMAP predicts that critical feature patterns to which humans learn to pay attention are stored in memory. Under language/cultural supervision, these prototypes can be either specific ("exemplars"; Estes 1994; Medin & Smith, 1981; Medin et al. 1983) or general ("prototypes"; Posner & Keele 1970; Smith & Minda 1998; 2000; Smith et al. 1997). Typically, both specific and general information will be learned ("rule-plus-exceptions"; Nosofsky 1984; 1987; Nosofsky et al. 1992; Palmeri et al. 1994).

ARTMAP is a standard tool for learning complex categorical relationships from high-dimensional input vectors that include color among other visual features, while autonomously discovering hierarchical knowledge relationships among the categories (Carpenter et al. 2004a; 2004b; Parsons & Carpenter 2003).

S&B summarize familiar features of neural models of supervised learning using nomenclature about games. Although these games sound novel, they actually embody well-known neural modeling concepts, including memory search or hypothesis testing to create new categories, the use of predictive success to culturally constrain learned naming, and the need to control category size. All of these properties are unified and proceed automatically in ARTMAP algorithms. It remains for S&B to demonstrate, through comparative benchmarks, that their models can cope with the categorical challenges that this alternative approach has already handled.

Modeling category coordination: Comments and complications

James A. Hampton
Psychology Department, City University, Northampton Square, London EC1V OHB, United Kingdom. **hampton@city.ac.uk**
www.staff.city.ac.uk/hampton

Abstract: Consideration of color alone can give a misleading impression of the three approaches to category coordination: the nativist, empiricist and culturalist models. Empiricist models can benefit from a wider range of correlational information in the environment. Also, all three approaches may explain a set of perceptual categories within the human repertoire. Finally, a suggestion is offered for supplementing the naming game by varying the social status of agents.

The broad conclusion drawn by Steels & Belpaeme (S&B) on the basis of their explorations of three general models for category name coordination is that, whereas the genetic and the cultural/language-based models can lead to coordinated categorization and naming practices within populations, the statistical structure available for colors in the immediate environment is insufficient to allow the empiricist model to achieve the same level of performance.

The latter claim is critically dependent on the choice of environmental structure provided to the empiricist model. S&B explored only one source of structure – a random sampling of pixels taken from photos of environmental scenes. Before drawing any firm conclusion about the possibility of achieving full coordination of categories simply from statistical covariation in the world, a more realistic characterization of the environment is surely necessary. In particular, colors are not seen by individuals as independent pixels, but as reflectances of the surfaces of objects and parts of the visual scene, which can be tracked through space and time as the individual moves through the scene. Other visual properties such as shape and size of the color patch, where it is located relative to objects in the scene, and sensory properties from other modalities all provide rich sources of correlational structure which establish a categorization of the world. Coordination of color categories may then benefit from the association of colors with object classes (oranges are typically orange, lemons typically yellow, the sky typically blue, blood red, and so forth). To take one of their examples, if coordination of color categories is important for the detection of poisonous mushrooms, then it is unrealistic to suppose that the morphology, size, smell, and habitat of the mushroom will not also play a role in categorization – and hence provide crucial evidence about where to draw the color category boundary in this instance. It is therefore an empirical question whether a richer modeling of the statistical structure in the environment would be sufficient to allow a purely empiricist model to develop coordinated categories as efficiently as the other two model types.

Presenting the three approaches to the problem as mutually contrasting accounts may also be misleading. Human categorizers (and human cultures that develop category systems) form and name categories on the basis of a wide range of sources of information. It is easy to find prima facie candidates for perceptual categories that are grounded in each of the three models explored – genetic, empirical, or cultural. Coordination of names for basic emotions such as happiness and grief, or bodily states such as hunger, thirst, or fatigue is presumably based on our common genetics. Coordination of names for biological classes probably relies on the fact that the similarity structure of biological classes at an intermediate level (e.g., elephant, tiger) contains clearly defined clusters with high within-cluster similarity and low between-cluster similarity, giving relatively universal taxonomic systems across different cultures at this level (Lopez et al. 1997). Artifact classification at the basic level may similarly rely solely on high levels of distinctiveness (Rosch et al. 1976). Other categories that depend more on language may be found in culturally-specific categories relating to social practices. For example, classification of ceramics, painting, or music in terms of different artistic styles, or notions of good and bad taste in clothing or decoration are perceptually grounded, but may depend heavily on language for their coordination. It is only the fact of having the concept in the language that leads the language learner to attend to the relevant perceptual cues and construct the necessary prototype representations. A wider view of perceptual categories suggests therefore that the three approaches considered by S&B – nativism, empiricism, and culturalism – all have their place in explaining the rich repertoire of human concepts.

My final comment relates to the cultural model itself. S&B's model assumes a fully cooperative pair of individuals in the language game. Each is willing to adapt his/her categorization and usage of language in the service of improving communication. In actual human societies, the degree of cooperation may be less evident. The right to determine how things are categorized and named is not so evenly distributed. Children are expected to conform to adult norms. In Western cultures, social classes whose education has given them greater access to the elaborated use of language (Bernstein 1981) may determine that there are right and wrong ways in which to classify and name the world. For an increasing number of domains, the correct use of a word is the province of an expert – Putnam's Division of Linguistic Labor (Putnam 1975) – to whom other language users are inclined (and expected) to defer (Kalish 1995). It would be an interesting exercise for S&B to consider introducing differing levels of social status in the naming game. Agents of lower status would be willing to adapt their representations rapidly, whereas agents of higher status would hold on to their beliefs longer in the face of disagreement, particularly if the interlocutor was of lower status than themselves. It would be fascinating to know if introducing this dimension into the game leads to quicker coordination of categories throughout the community, and whether the higher or lower status players end up showing greater or less variance as a subclass of agents.

Language and the game of life

Stevan Harnad

Centre de Neuroscience de la Cognition, Université du Québec à Montréal, Montréal, Québec H3C 3P8, Canada. **harnad@uqam.ca**
http://www.ecs.soton.ac.uk/~harnad/

Abstract: Steels & Belpaeme's (S&B's) simulations contain all the right components, but they are put together wrongly. Color categories are unrepresentative of categories in general and language is not merely naming. Language evolved because it provided a powerful new way to acquire categories (through instruction, rather than just the old way of other species, through trial-and-error experience). It did not evolve so that multiple agents looking at the same objects could let one another know which of the objects they had in mind, co-coining names for them on the fly.

Contra Wittgenstein (1953), language is not a game. (Maynard-Smith [1982] would no doubt plead *nolo contendere*.) The game is *life*, and language evolved (and continues to perform) in life's service – although it has since gained a certain measure of autonomy too.

So are Steels & Belpaeme (S&B) inquiring into the functional role for which language evolved, the supplementary roles for which it has since been coopted, or merely the role something possibly resembling language might play in robotics (another supplement to our lives)?

For if S&B are studying the functional role for which language evolved, that role is almost certainly absent from the experimental conditions that they are simulating. Their computer simulations do not capture the ecological conditions under which, and for which, language began. The tasks and environments set for S&B's simulated creatures were not those that faced any human or prehuman ancestor, nor would they have led to the evolution of language had they been. On the contrary, the tasks faced by our prelinguistic ancestors (as well as our nonlinguistic contemporaries) are precisely the ones *left out* of S&B's simulations.

S&B do make two fleeting references to a world in which foragers need to learn to recognize and sort mushrooms by *kind* – with color possibly serving as one of the features on the basis of which they sort. But a task like learning to sort mushrooms by kind is not what S&B simulate here. They simulate the task of sorting colors, and not by *kind*, but by a kind of "odd man out" exercise called the "discrimination game." In this game, the agent sees a number of different colors (the "context"), of which one (the "topic") is the one that must be discriminated from the rest. If this is done by two agents, it is called the "guessing game," with the speaker both discriminating and naming the topic-color, and the hearer having to guess which of the visible context-colors the speaker named. Both agents see all the context-colors.

Now the first thing we must ask is: (i) Were any of our prelinguistic ancestors ever faced with a task anything like this? (ii) And if they had been, would that have led to the evolution of language? (iii) Indeed, is what is going on in S&B's task *language* at all?

I would like to suggest that the answer to all three questions is no. S&B's is not an ecologically valid task; it is not a canonical problem that our prelinguistic ancestors encountered, for which language evolved as the solution. And even if we trained contemporary animals to do something like it (as some comparative psychologists have done, e.g., Leavens et al. 1996), it would not be a linguistic task – indeed it would hardly even be a categorization task, but more like a joint multiple-choice task requiring some "mind-reading" (Premack & Woodruff 1978; Tomasello 1999) plus some coordination (Fussell & Krauss 1992; Markman & Makin 1998).

On the other hand, there is no doubt that our own ancestors, once language had evolved, *did* face tasks like this, and that language helped them perform such tasks. But language helps us perform many tasks (even learning to ride a bicycle or to do synchronized swimming) for which language is not necessary, for which it did not evolve, and which are not themselves linguistic tasks. (This is S&B's "chicken/egg" problem, but in a slightly different key.)

Let's now turn to something that *is* ecologically valid. Our prelinguistic ancestors (and their nonlinguistic contemporaries, as well as our own) did face the problem of categorization and category learning. They did have to know or learn what to do with different *kinds* of things, in order to survive and reproduce: what to eat or not eat, what to approach or avoid, what kind of thing to do with what kind of thing, and so forth. But categorizing is not the same as discriminating (Harnad 1987). We discriminate things that are present simultaneously, or in close succession; hence, discrimination is a *relative* judgment, not an *absolute* one. You don't have to *identify* what the things are in order to be able to discern whether two things are the same thing or different things, or whether this thing is more like that thing or that thing. Categorization, in contrast, calls for an absolute judgment of a thing in isolation: "What kind of thing is this?" And the identification need not be a name; it can simply be *doing* the kind of thing that you need to do with that kind of thing (flee from it, mate with it, or gather and save it for a rainy day).

So categorization tasks have not only ecological validity, but cognitive universality (Harnad 2004). None of our fancier cognitive capacities would be possible if we could not categorize. In particular, if we could not categorize, we could not name. To be able to identify a thing correctly, in isolation, with its name, I need to be able to discriminate it absolutely, not just relatively – that is, not just from the alternatives that happen to be copresent with it at the time (S&B's "context"), but from all other things I encounter, past, present, and (one hopes) future, with which it could be *confused*. (Categorization is not necessarily exact and infallible. I may be able to name things correctly based on what I have sampled to date, but tomorrow I may encounter an example that I not only cannot categorize correctly, but that shows that all my categorization to date has been merely approximate too.)

Notice that I said categorize *correctly*. That is the other element missing from S&B's analyses. For S&B, there are three ways in which things can be categorized: (N) innately ("nativism"), (E) experientially ("empiricism"), and (C) culturally ("culturalism," although one wonders why S&B consider cultural effects nonempirical!). To be fair, the way S&B put it is that these are the three ways in which categories can come to be *shared* – but clearly one must *have* categories before one can share them (the chicken/egg problem again!).

Where do the S&B agents' color categories come from? They seem to think that categories come from the "statistical structure" of the things in the world, such as how much things resemble one another physically, how frequently they occur and cooccur, and how this is reflected in their effects on our sensorimotor transducers. This is the gist of S&B's factor E, empiricism. Where the statistical structure has been picked up by evolution (another em-